Reference Architecture

Version: 0.2.0

Type: Draft

Release Date: August 15, 2024

Canada Health Infoway

Version	0.2.0
Type	Draft Pre-Ballot
Release Date	August 15, 2024

This version of the Reference Architecture is referenced by and/or references the following Interoperability products:

- pan-Canadian Patient Summary Interoperability Specifications (PS-CA v2.0.0 DFT-preBallot)
- pan-Canadian FHIR Exchange Interoperability Specifications (CA:FeX v2.1.0 DFT-preBallot)

1 Reference Architecture (RA)

Version	0.2.0
Type	Draft Pre-Ballot
Release Date	August 15, 2024

This version of the Reference Architecture is referenced by and/or references the following Interoperability products:

- pan-Canadian Patient Summary Interoperability Specifications (PS-CA v2.0.0 DFT-preBallot)
- pan-Canadian FHIR Exchange Interoperability Specifications (CA:FeX v2.1.0 DFT-preBallot)

1.1 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.

Term / Acronym	Meaning
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.
CA:FMT	Canadian Formatting Service (CA:FMT) is a Canadian Integration Specification that provides formatting support service. It provides support for transformation of documents between different formats (e.g. from FHIR to PDF, CDA, etc.).
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.
Clinical Data Repository (Local or Central)	A Clinical Data Repository (i.e., document 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.
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. PS-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.

Term / Acronym	Meaning
CT	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. (Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-7.html)
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.
Document Repository (Local or Central)	See Clinical Data Repository (Local or Central)
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)	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: • 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	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. *Note: Extensible PS-CA Dataset refers to the addition of data domains such as Family History.
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®)	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. (Source: http://www.hl7.org/implement/standards/fhir/)
Information/ Semantic Interoperability Requirements	Requirements for syntax and semantics such that data exchanged between health record systems can be interpreted and the meaning of the data ascertained.
Integrating the Healthcare Enterprise (IHE)	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. (Source: https://www.ihe.net/)
IHE Actor	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.). (Source: https://wiki.ihe.net/index.php/Actors)
IHE Domain	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). (Source: https://wiki.ihe.net/index.php/Domains)
IHE Integration Profiles	IHE Integration Profiles provide a solution to the interoperability challenges which have arisen in daily clinical work, as described in the Use Cases. Integration Profiles include detailed technical specifications for the use and implementation of relevant standards thus ensuring an uninterrupted flow of information between different healthcare IT applications in support of the specific use case. The Profiles describe how healthcare IT systems can provide integrated support for a clearly defined workflow, each of which individually supports a clinical task within a specific clinical domain. IHE Profiles can be used for a step-by-step implementation of systems in different domains and the gradual building of interoperable eHealth applications. (Source: https://www.ihe-europe.net/about-us/faq)

Term / Acronym	Meaning
IHE Transactions	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.). (Source: https://wiki.ihe.net/index.php/PCC_TF-1/About)
International Patient Summary (IPS)	The IPS is 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. (Source: https://wiki.ihe.net/index.php/International_Patient_Summary_(IPS))
Interoperability	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. 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. For more information about interoperability, please visit Canada Health Infoway - Interoperability.
IUA	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. (Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-34.html)
Local, Scheduled Care	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.
Local, Unscheduled Care	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.
Longitudinal Electronic Health Record	A longitudinal electronic health record is a single comprehensive patient record comprised of data from numerous data sources across the healthcare continuum.

Term / Acronym	Meaning
Medical Home	The College of Family Physicians of Canada describes the Medical Home as: "The Patient's Medical Home (PMH) is a family practice defined by its patients as the place they feel most comfortable—most at home—to present and discuss their personal and family health and medical concerns. It is the central hub for the timely provision and coordination of a comprehensive menu of health and medical services patients need." To read more about the Patient's Medical Home, please visit The College of Family Physicians of Canada's published document, A Vision for Canada - Family Practice - The Patient's Medical Home.
MHD	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. (Source: https://profiles.ihe.net/ITI/MHD/index.html)
On-Demand	Refers to the capability to generate a patient summary at the time it is requested. This means retrieving a patient's most current health data from available sources (e.g., CDR, EHR) when needed, ensuring timely access to information for clinical decision-making and patient care.
Patient Portal	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.
Patient Proxy	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.
Patient Summary-CA (PS-CA)	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.
PDQm	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. (Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-38.html)

Term / Acronym	Meaning
PIXm	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. (Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-41.html)
PMIR	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. (Source: https://www.ihe.net/uploadedFiles/Documents/ITI/IHE_ITI_Suppl_PMIR.pdf)
Producer (e.g., PS-CA Producer)	A health records system (e.g., EMR, HIS, CIS, PHR, or EHR) that creates/produces a clinical document (e.g. PS-CA) in response to a request from an authorized health care provider, the subject of care or another authorized health records system.
Projectathon	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.
PS-CA Solution	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.
PS-CA Specifications	**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.
PT	Provinces and Territories

Term / Acronym	Meaning
RA	The Reference Architecture (RA) is intended as an evolving blueprint of service availability that supports a broader interoperability landscape, not limited to patient summaries. Its purpose is to facilitate multi-stakeholder dialogue, collaboration and convergence towards common, open standards. It is a conceptual technical view that provides a common vocabulary and a set of actors and transactions representing typical components in a digital health ecosystem (public and private sector solutions). It is combination of building blocks adopted from international standards development bodies and Canadian developed implementation patterns.
Shareable Health Link (SHL, SHLink)	A secure, standardized link that enables patients to share their clinical health information (e.g., patient summaries) with healthcare providers. SHLs can be shared via QR codes or other secure electronic methods (e.g., email). SHLs are designed to facilitate the smooth exchange of health data, allowing information to travel with patients wherever they might be, supporting the continuity of care. SHL is based on HL7 SMART Health Links and generally can be used interchangeably.
SUT	System Under Test
SVCM	Sharing Valuesets, Codes and Maps (SVCM) defines a lightweight interface through which healthcare systems may retrieve centrally managed uniform nomenclature and mappings between code systems based on the HL7 Fast Healthcare Interoperability Resources (FHIR) specification. (Source: https://wiki.ihe.net/index.php/Sharing_Valuesets,_Codes_and_Maps_(SVCM))
Technical Interoperability Requirements	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.
Terminology	Collection of uniquely identifiable concepts with associated representations, designations, associations and meanings.
XDM	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. (Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-16.html)
XDS	The Cross-Enterprise Document Sharing (XDS) IHE Integration Profile facilitates the registration, distribution and access across health enterprises of patient electronic health records. (Source: https://profiles.ihe.net/ITI/TF/Volume1/ch-10.html)

1.2 Overview

1.2.1 Purpose

The Reference Architecture (RA) content, to support interoperability, provides guidance on how to apply specific integration patterns that improve the way computer systems in healthcare share information in Canada. Systems that implement these capabilities can be represented by standardized interoperability patterns that are reusable across a wider array of eco-system solutions. The RA focuses on highlighting key elements of standardization that can be employed in solution design at system level. The elements presented in here are patterns that will be referred to in Companion Guides and Interoperability Specification documentation of specific priority areas (e.g., Patient Summary, e-Referral, etc.).

A large number of the integration patterns presented herein are interoperability profiles developed by IHE International. These patterns are based on open standards and are regularly tested at international interoperability testing events - Connectathons. This provides great value to both vendors and purchasers in that it creates a catalog of capabilities that can be used for system wide eco-system design. By following best practices and international standards, this Reference Architecture can grow over time to represent a pan-Canadian point of view and become an approach to interoperability.

In addition to existing published IHE profiles, the Reference Architecture also provides detailed implementation guidance for areas that contain Canadian specific implementation requirements (e.g., how to use IDP profiles - IUA, guidance on security and auditing - ATNA, etc.).

1.2.2 Intended Audience

The intended audience of the 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 vendors participating in the IHE initiative
- 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.

1.2.3 The Reference Architecture Overview

The Reference Architecture contains information relevant to system level integration patterns, supporting interoperability. Other interoperability specifications (e.g., PS-CA, CA:FeX) may contain project-specific references to architecture and/or refer directly to this Reference Architecture document.

Program-level: The RA is a living document that describes IHE profiles that have been recommended to be adopted by jurisdictions and vendors to improve the way computer systems in healthcare share information in Canada. Additional IHE profiles will be added over the long-term as new priorities are included. Program level information of the Reference Architecture can be found in this space and page (see diagram below), where summarized information for the IHE profiles have been included and grouped under the following categories:

- Foundational Profiles: Includes a list of foundational IHE profiles that perform basic functions within the ecosystem such as authorized log-in, audit event logging, consistent time synchronization, etc.
- Document/Data Sharing Profiles: Includes a list of IHE profiles that provide recommended standards for document and data sharing between clinical solutions.
- Patient Identity Profiles: Includes a list of IHE profiles that enable confirming the identity of the patient or subject of care.

Project-specific: In addition to the Reference Architecture, there may also be project-specific references to architecture that contain a subset of the recommended profiles from this document. These project specific guides will list sequence diagrams that support the needs of those projects. Actors and transactions can be grouped together from multiple profiles to address the business requirements of the project specific use cases.

1.2.4 How to Read the Reference Architecture

- The Reference Architecture lists the actors and transactions from the IHE profiles that can enable secure exchange of health information in Canada.
- The interactions and standardized transactions between the actors as defined by the IHE methodology framework 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, however, these latter roles can also be assumed by vendor systems, the RA does not prevent this.
- Within the Reference Architecture, options for implementation have been highlighted, with Option 1 having two scenarios.
    - Option 1, Scenario #1: MHD implementation, where the Document Repository is Central
    - Option 1, Scenario #2: MHD implementation, where the Document Repository is Local
    - Option 2: CA:FeX (Canadian FHIR Exchange) Implementation
- A preferred option is indicated with an asterisk *(e.g. Option 1, Scenario #1)
- A list of all the IHE profiles are included at the bottom of the Reference Architecture
- A list of Canadian National Integration Specifications and Guidance are listed at the bottom of the Reference Architecture, where national extensions to existing IHE profiles and net new profiles for Canada are labelled distinctly.
- A legend is found at the bottom of the Reference Architecture to inform readers about the details of the diagram.

1.2.5 How to Use the Reference Architecture

The list below summarizes how to use this Reference Architecture 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, 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 Reference Architecture is a living document and will evolve based on feedback from all stakeholders and further refinements driven by implementation specific uses cases.

This document is published 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 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: The reader is expected to have a moderate degree of familiarity with IHE profiles, especially the ones listed in the Reference Architecture Overview diagram (e.g. IUA, ATNA, CT, MHD).

1.2.6 IHE Profile Versions

The following describes the published versions in scope for the IHE Profiles that have been referenced in this Reference Architecture.

- MHD: v4.1.0: Trial Implementation (2022-03-01), based on FHIR R4
- IUA: Revision 2.2 – Trial Implementation (2022-06-17)
- ATNA: Revision 19.0 – Final Text (2022-06-17)
- RESTful ATNA Supplement: Rev. 3.3 – Trial Implementation (2021-07-02), based on FHIR R4
- CT: Revision 19.0 – Final Text (2022-06-17)

For any other IHE profiles that are referenced in the Reference Architecture, but not mentioned in the above list, the latest version is considered.

1.2.7 High-Level View of the Reference Architecture

This is a high-level view of relevant IHE Profiles and Interoperability Specifications to support the Reference Architecture. The view contains a superset of profiles that offer alternatives to exchanging documents and data, depending on Jurisdictional service type and availability.

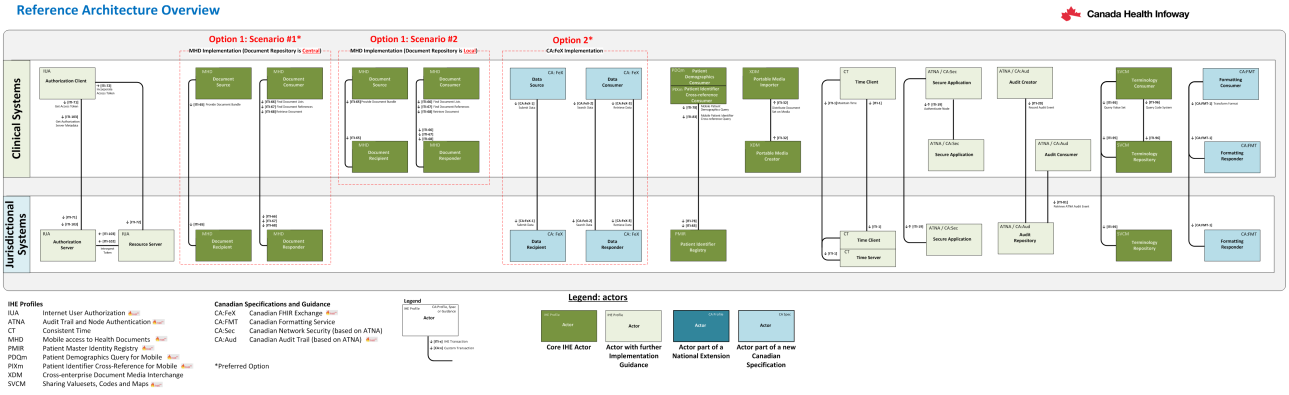

Reference Architecture Overview
Canada Health Infoway
Option 1: Scenario #1*
MHD Implementation (Document Repository is Central)
Option 1: Scenario #2
MHD Implementation (Document Repository is Local)
Option 2*
CA:FeX Implementation
Clinical Systems
Jurisdictional Systems
IUA Authorization Client
[ITI-72] Incorporate Access Token
[ITI-71] Get Access Token
[ITI-103] Get Authorization Server Metadata
IUA Authorization Server
IUA Resource Server
[ITI-103]
[ITI-102] Introspect Token
MHD Document Source
MHD Document Consumer
[ITI-65] Provide Document Bundle
[ITI-66] Find Document Lists
[ITI-67] Find Document References
[ITI-68] Retrieve Document
MHD Document Recipient
MHD Document Responder
CA: FeX Data Source
CA: FeX Data Consumer
[CA:FeX-1] Submit Data
[CA:FeX-2] Search Data
[CA:FeX-3] Retrieve Data
CA: FeX Data Recipient
CA: FeX Data Responder
PDQm Patient Demographics Consumer
PIXm Patient Identifier Cross-reference Consumer
[ITI-78] Mobile Patient Demographics Query
[ITI-83] Mobile Patient Identifier Cross-reference Query
PMIR Patient Identifier Registry
XDM Portable Media Importer
[ITI-32] Distribute Document Set on Media
XDM Portable Media Creator
CT Time Client
[ITI-1] Maintain Time
CT Time Server
ATNA / CA:Sec Secure Application
[ITI-19] Authenticate Node
ATNA / CA:Aud Audit Creator
[ITI-20] Record Audit Event
ATNA / CA:Aud Audit Consumer
[ITI-81] Retrieve ATNA Audit Event
ATNA / CA:Aud Audit Repository
SVCM Terminology Consumer
[ITI-95] Query Value Set
[ITI-96] Query Code System
SVCM Terminology Repository
CA:FMT Formatting Consumer
[CA:FMT-1] Transform Format
CA:FMT Formatting Responder
IHE Profiles
IUA Internet User Authorization
ATNA Audit Trail and Node Authentication
CT Consistent Time
MHD Mobile access to Health Documents
PMIR Patient Master Identity Registry
PDQm Patient Demographics Query for Mobile
PIXm Patient Identifier Cross-Reference for Mobile
XDM Cross-enterprise Document Media Interchange
SVCM Sharing Valuesets, Codes and Maps
Canadian Specifications and Guidance
CA:FeX Canadian FHIR Exchange
CA:FMT Canadian Formatting Service
CA:Sec Canadian Network Security (based on ATNA)
CA:Aud Canadian Audit Trail (based on ATNA)
*Preferred Option
Legend
IHE Profile
CA:Profile, Spec or Guidance
Actor
[ITI-x] IHE Transaction
[CA:x] Custom Transaction
Legend: actors
IHE Profile Actor
Core IHE Actor
IHE Profile Actor
Actor with further Implementation Guidance
CA Profile Actor
Actor part of a National Extension
CA Spec Actor
Actor part of a new Canadian Specification

1.3 Foundational Profiles

Background

Foundational IHE Profiles and pan-Canadian Interoperability Specifications address critical interoperability issues such as user authorization (e.g. IUA), security node and audit records (e.g. ATNA), consistent time (e.g. CT), terminology (e.g. SVCM), document transformation/formatting (e.g. CA:FMT) and more.

Assumption

Vendors and jurisdictions in the ecosystem can optionally choose to play the standardized actors and transactions listed in the Foundational Profiles for the PS-CA Specification. 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 being addressed using solutions that currently exist in their respective enterprise architecture.

IHE Profiles, Specifications and Guidance included:

- Audit Trail and Node Authentication (ATNA)
    - Canadian Network Security (CA:Sec) Implementation Guidance
    - Canadian Audit Trail (CA:Aud) Implementation Guidance
- Internet User Authorization Profile (IUA)
- Consistent Time (CT)
- Sharing Valuesets, Codes and Maps (SVCM)
- Canadian Formatting Service (CA:FMT)

Legend

The below diagram is the legend for the sequence diagrams included for each profile, specification and guidance.

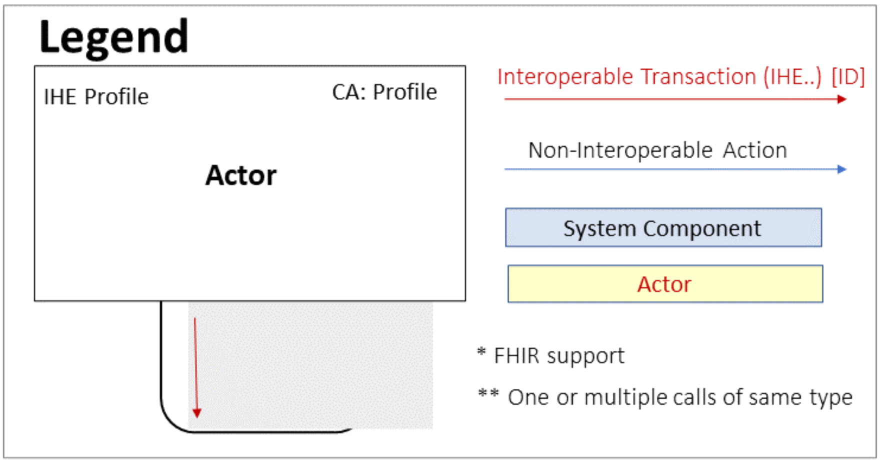

1.3.1 Audit Trail and Node Authentication (ATNA)

Overview

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.

For details, see IHE Audit Trail and Node Authentication (ATNA) profile and RESTful ATNA Supplement.

Actors and Transactions

The following diagram provides an overview of the ATNA profile Actors, Transactions and their interactions.

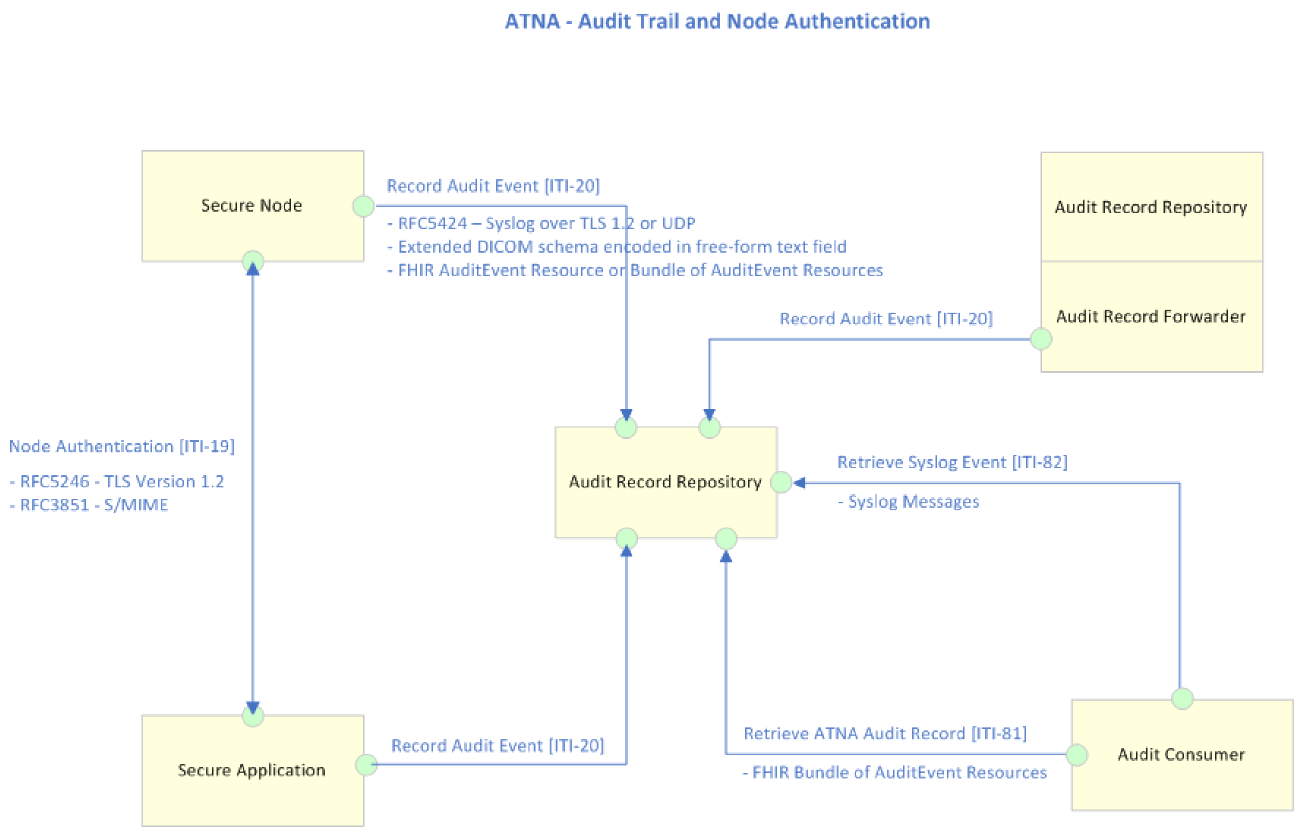
ATNA - Audit Trail and Node Authentication
Secure Node
Record Audit Event [ITI-20]
- RFC5424 – Syslog over TLS 1.2 or UDP
- Extended DICOM schema encoded in free-form text field
- FHIR AuditEvent Resource or Bundle of AuditEvent Resources
Audit Record Repository
Audit Record Forwarder
Record Audit Event [ITI-20]
Node Authentication [ITI-19]
- RFC5246 - TLS Version 1.2
- RFC3851 - S/MIME
Audit Record Repository
Retrieve Syslog Event [ITI-82]
- Syslog Messages
Secure Application
Record Audit Event [ITI-20]
Retrieve ATNA Audit Record [ITI-81]
- FHIR Bundle of AuditEvent Resources
Audit Consumer

The table below lists the transactions for each actor directly involved in the ATNA profile. To claim compliance with ATNA, an actor shall support all required transactions (labeled "R") and may support the optional transactions (labeled "O").

Actor	Transaction	Optionality
Audit Record Repository	Record Audit Event [ITI-20] Retrieve ATNA Audit Event [ITI-81] Retrieve Syslog Event [ITI-82]	R O O
Audit Consumer	Retrieve ATNA Audit Event [ITI-81] Retrieve Syslog Event [ITI-82]	O O
Audit Record Forwarder	Record Audit Event [ITI-20]	R

Secure Node	Authenticate Node [ITI-19] Record Audit Event [ITI-20]	R R
Secure Application	Authenticate Node [ITI-19] Record Audit Event [ITI-20]	R R

Transactions

- Authenticate Node [ITI-19] – In the Authenticate Node transaction, the local Secure Node presents its identity to a remote Secure Node and authenticates the identity of the remote node. After this mutual authentication, other secure transactions may take place through this secure pipe between the two nodes. Uses RFC5246 - Transport Layer Security (TLS) Protocol Version 1.2 and RFC3851 - Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification protocols.
- Record Audit Event [ITI-20] – This transaction is used to report auditable events to an Audit Record Repository. Uses RFC5424 – Syslog over TLS 1.2 or UDP protocols.
- Retrieve ATNA Audit Event [ITI-81] – This transaction is used to search ATNA events recorded in an ATNA Audit Record Repository. The result is a FHIR bundle of AuditEvent Resources that match a set of search parameters.
- Retrieve Syslog Event [ITI-82] – This transaction is used to retrieve syslog messages from the Audit Record Repository subject to parameters that limit the retrieval.

Sequence Diagram

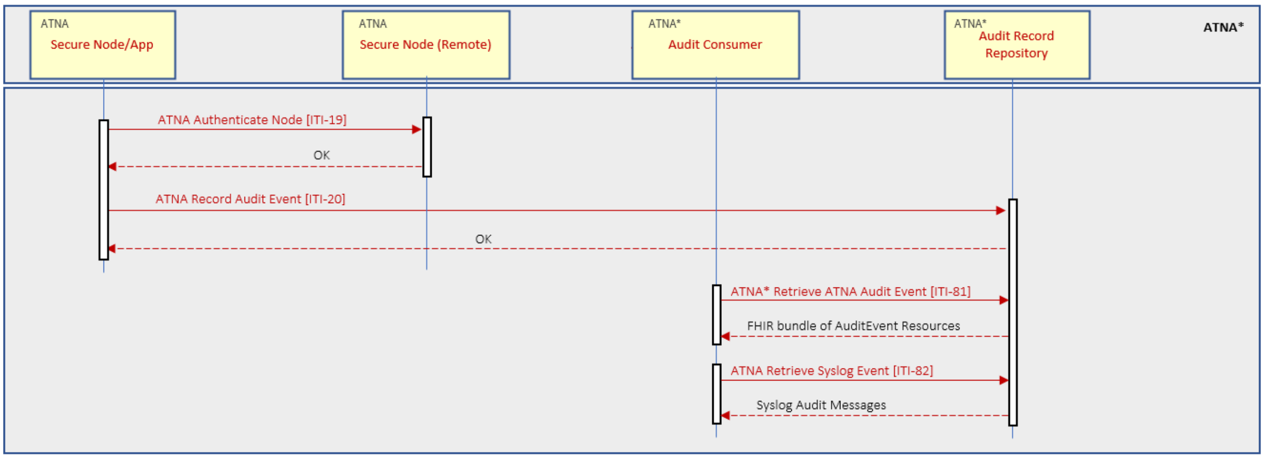
ATNA
Secure Node/App
ATNA
Secure Node (Remote)
ATNA*
Audit Consumer
ATNA*
Audit Record Repository
ATNA*
ATNA Authenticate Node [ITI-19]
OK
ATNA Record Audit Event [ITI-20]
OK
ATNA* Retrieve ATNA Audit Event [ITI-81]
FHIR bundle of AuditEvent Resources
ATNA Retrieve Syslog Event [ITI-82]
Syslog Audit Messages

Canadian Implementation Guidance for ATNA - CA:Sec and CA:Aud

The ATNA profile addresses two main concerns: Security and Event Logging for the purpose of Auditing. Given the fact that Security and Auditing are tightly coupled, along with the multiple options offered for both aspects, ATNA is a complex profile with extensive documentation.

CA:Sec and CA:Aud implementation guidance were introduced to allow for a lightweight ATNA, bring improvements by decoupling the two main aspects of ATNA: Security and Audit, and focus on a few options for modern formats and technologies. This guidance is not replacing ATNA. An implementation that is already compliant with ATNA will be able to pass ATNA tests.

The following diagram presents how the Canadian implementation guidance has segmented the key components of ATNA.

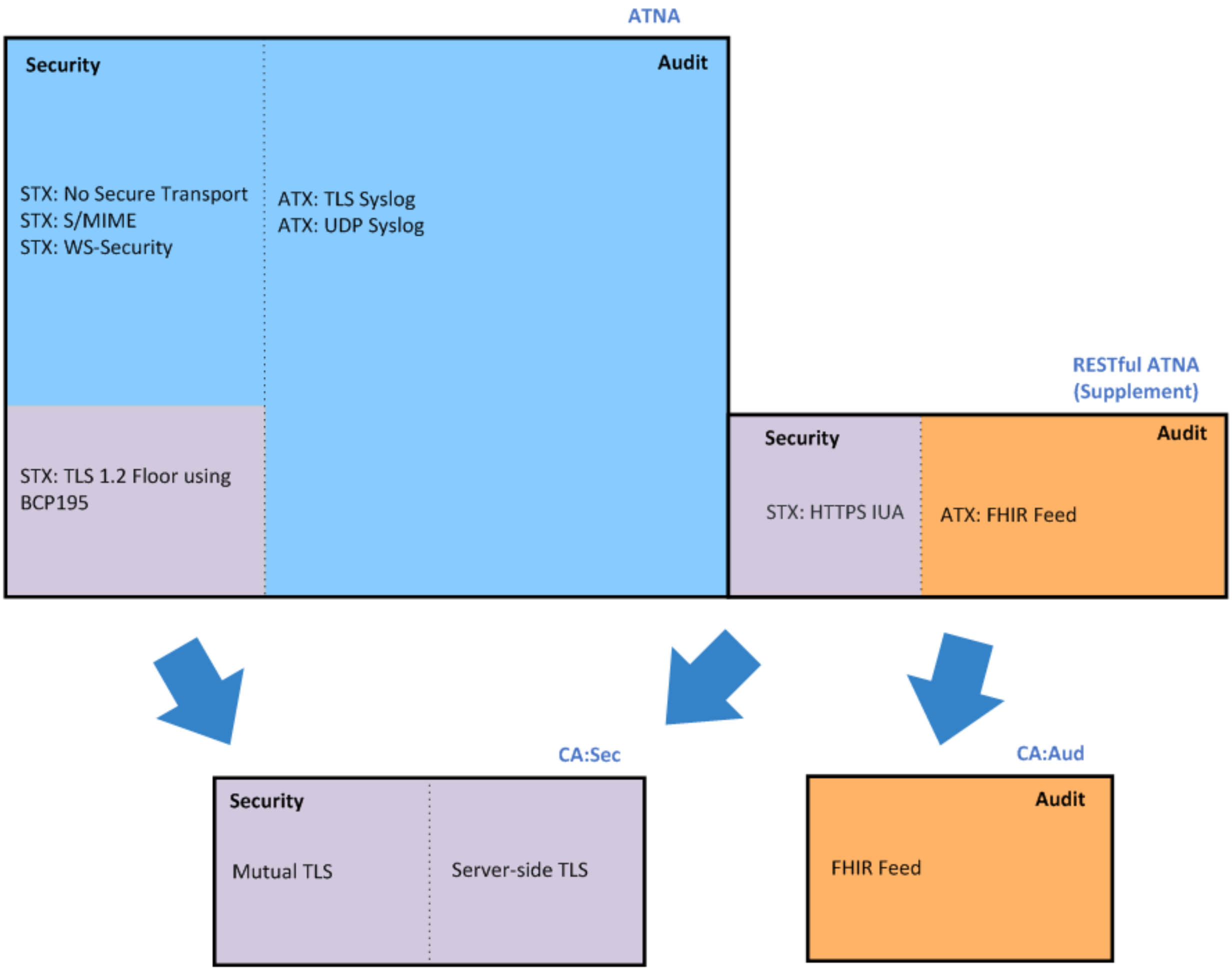

The section below provides comparison tables between the full ATNA profile and the options selected for CA:Sec and CA:Aud.

The following notation definitions are used throughout this section:

Optionality notation is defined as:

R	Required
O	Optional

Transport Protocol notation is defined as:

STX prefix	Secure transport protocol
ATX prefix	Audit transport protocol

Actors/Transactions

ATNA

Actors	Transactions	Optionality
Secure Node	Authenticate Node [ITI-19]	R
	Record Audit Event [ITI-20]	R
Secure Application	Authenticate Node [ITI-19]	R
	Record Audit Event [ITI-20]	R
Audit Record Repository	Record Audit Event [ITI-20]	R
	Retrieve ATNA Audit Event [ITI-81]	O
	Retrieve Syslog Event [ITI-82]	O
Audit Consumer	Retrieve ATNA Audit Event [ITI-81]	O
	Retrieve Syslog Event [ITI-82]	O
Audit Record Forwarder	Record Audit Event [ITI-20]	R

CA:Sec

Actors	Transactions	Optionality
Secure Application	Authenticate Node [ITI-19]	R

CA:Aud

Actors	Transactions	Optionality
Audit Creator	Record Audit Event [ITI-20]	O (Note 1)

Audit Record Repository	Record Audit Event [ITI-20]	O (Note 1)
	Retrieve ATNA Audit Event [ITI-81]	O (Note 2)
Audit Record Forwarder	Record Audit Event [ITI-20]	O (Note 1)
Audit Consumer	Retrieve ATNA Audit Event [ITI-81]	R

Note 1: The audit events must be recorded using the IHE Record Audit Event [ITI-20] with FHIR Feed option or other (IHE or non-IHE) methods.

Note 2: This transaction is required if the Audit Record Repository is central.

Notes

- ATNA defines two actors with similar role: Secure Node and Secure Application, that fulfill roles in both Security and Audit aspects via mandatory transaction requirements. This causes Security and Audit to be tightly coupled, meaning that to exchange secure communication, auditing also must be implemented as defined by ATNA.
  The audit messages must be recorded by means defined by ATNA.
- CA:Sec defines a single actor and a single transaction for secure communication.
- CA:Aud defines actors that are responsible for auditing only. Secure communication is recommended to be achieved via actor grouping with the CA:Sec actor.
  The audit messages can be recorded by any means, using either IHE transaction ITI-20 with FHIR option or any other (IHE or non-IHE) methods. The messages must be made available for retrieval in FHIR format via IHE transaction ITI-81.

Actor Options

ATNA

Actor	Options
Audit Record Repository	Retrieve Audit Message
	Retrieve Syslog Message
	ATX: FHIR Feed
	ATX: TLS Syslog
	ATX: UDP Syslog
Audit Consumer	Retrieve Audit Message

	Retrieve Syslog Message
Audit Record Forwarder	ATX: FHIR Feed
	ATX: TLS Syslog
	ATX: UDP Syslog
Secure Node	Radiology Audit Trail
	FQDN Validation of Server Certificate
	STX: No Secure Transport
	STX: TLS 1.2 Floor using BCP195
	STX: S/MIME
	STX: WS-Security
	STX: HTTPS IUA
	ATX: FHIR Feed
	ATX: TLS Syslog
	ATX: UDP Syslog
Secure Application	Radiology Audit Trail
	FQDN Validation of Server Certificate
	STX: No Secure Transport
	STX: TLS 1.2 Floor using BCP195
	STX: S/MIME
	STX: WS-Security

	STX: HTTPS IUA
	ATX: FHIR Feed
	ATX: TLS Syslog
	ATX: UDP Syslog

CA:Sec

Actor	Options	Optionality
Secure Application	Mutual TLS	O (Note 1)
	Server-side TLS	O (Note 1)
	FQDN Validation of Server Certificate	R

Note 1: The Secure Application shall support one of the following options: Mutual TLS or Server-side TLS.

CA:Aud

Actor	Options	Optionality
Audit Creator	FHIR Feed	O (Note 1)
Audit Record Repository	FHIR Feed	O (Note 1)
	Retrieve Audit Message	O (Note 2)
Audit Record Forwarder	FHIR Feed	O (Note 1)
Audit Consumer	Retrieve Audit Message	R

Note 1: The audit events must be recorded using the IHE Record Audit Event [ITI-20] with FHIR Feed option or other (IHE or non-IHE) methods.

Note 2: This transaction is required if the Audit Record Repository is central.

Notes:

- ATNA offers many actor options
- CA:Sec and CA:Aud are focusing on a small subset of the ATNA options.
- CA:Sec options also improve security with recommendations for higher versions of TLS protocol and stronger cipher suites

Required Actor Groupings

ATNA

ATNA Actor	Actor(s) to be grouped with
Audit Record Repository	Consistent Time / Time Client
	ATNA / Secure Node or Secure Application
Audit Consumer	ATNA / Secure Node or Secure Application
Audit Record Forwarder	Consistent Time / Time Client
	ATNA / Secure Node or Secure Application
	ATNA / Audit Record Repository
Secure Node	Consistent Time / Time Client
Secure Application	Consistent Time / Time Client

CA:Sec and CA:Aud

None

Notes:

- ATNA requires multiple mandatory actor groupings.
- CA:Sec and CA:Aud do not require mandatory actor groupings.

Actor grouping is optional and is recommended to achieve additional functionality such as System Time Synchronization or Security and Auditing.

Canadian Network Security (CA:Sec) Implementation Guidance

The **CA:Sec (Canadian Network Security)** Implementation Guidance specifies the foundational elements needed to securely execute transactions between two systems.

CA:Sec is based on the ATNA profile and aims to bring improvements via loose coupling, and high cohesion, with focus on secure communication.

For more details see IHE ATNA profile, RESTful ATNA Supplement and ATNA ITI-19 transaction documentation.

Note: In addition to secure communication that is covered in CA:Sec, there are other critical aspects that need to be implemented to achieve a high degree of cybersecurity, see Security Considerations.

CA:Sec Description

The purpose is to define the actors directly involved in the CA:Sec Implementation Guidance, the transactions between them, and the options and groupings of the actors with other IHE profiles.

CA:Sec Actors/Transactions

The diagram below shows the actors directly involved in CA:Sec and the relevant transactions between them.

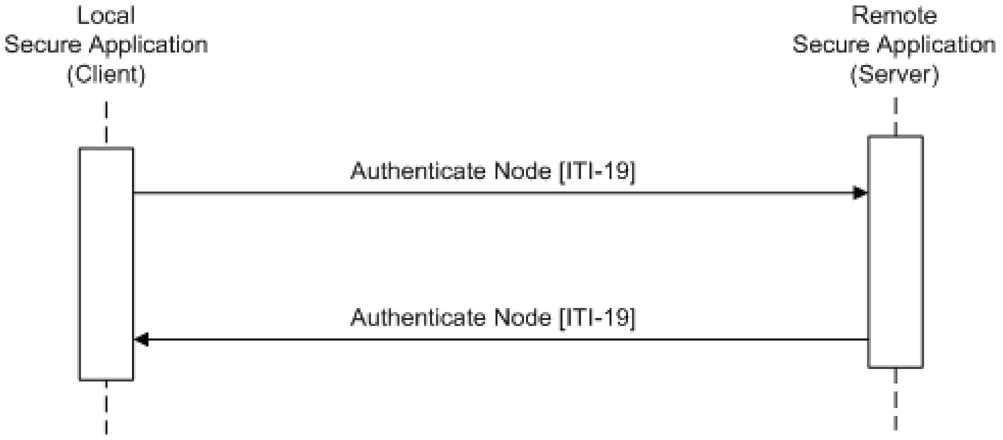

Mutual authentication (both client and server authenticate themselves)

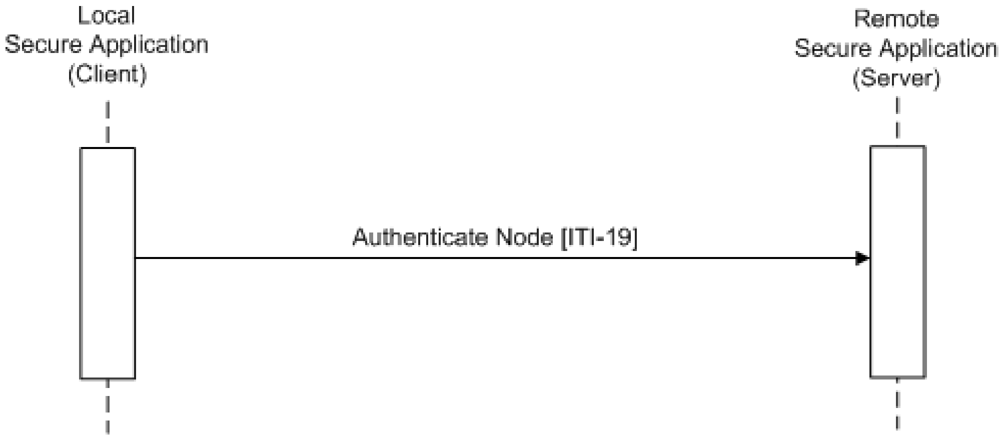

Server-side authentication (only the server authenticates itself)

The table below lists the transactions for each actor directly involved in CA:Sec.

To claim compliance with CA:Sec, an actor shall support all required transactions (labeled "R")

Actors	**Transactions**	**Optionality**
Secure Application	Authenticate Node [ITI-19]	R

Actor Descriptions and Requirements

Secure Application

A Secure Application provides security and privacy services (user authentication, secure communications, security audit recording, and security policy enforcement) for both grouped IHE actors and for functionality provided by related software and services within control of the Secure Application. Generally, the responsibilities of a Secure Application do not include the security of its environment, e.g., the operating system and database outside of its control. A smartphone app is an example of a Secure Application that has control over the security for the application, but not for the rest of the mobile device software or hardware.

Note that a Secure Application actor is inclusive of clients or servers, thus a Secure Application actor can represent a server agent running on a cloud platform.

A special case of a Secure Application is a Secure Node, that has complete control over the full stack from hardware to user interface and external communications. An ultrasound machine is an example of a Secure Node.

The Secure Application shall:

1. Use the Authenticate Node transaction for all network connections to or from the application that may expose private information as specified in Authenticate Node [ITI-19].
2. Provide sufficient authentication methods to ensure that only authorized users access the Secure Application.

CA:Sec Actor Options

For each actor in CA:Sec, the options labeled “R” shall be selected.

Actor	**Options**	**Optionality**
Secure Application	Mutual TLS	O (Note 1)
	Server-side TLS	O (Note 1)
	FQDN Validation of Server Certificate	R

Note 1: The Secure Application shall support one of the following options: Mutual TLS or Server-side TLS.

Mutual TLS Option

Mutual TLS option is a two-way authentication mechanism where both the server and the client applications identify themselves via the Transport Layer Security (TLS) protocol with the use of digital certificates.

Actors that support this option have the ability to both:

- Operate with the highest level of cyber protection for the TLS-protected communication channel per the current standards and IETF Best Current Practice (BCP195, RFC5246 at the time of writing – with TLS 1.2 or higher and selected cipher suites), and
- Restrict to the use of the current version of TLS (1.2 at the time of writing) [RFC5246] or higher, with strong recommendation for support of TLS version 1.3 [RFC8446].

Note: The recommendation for support of higher versions of TLS (1.3 at the time of writing) will become mandatory in the future.

An actor that supports this option shall be able to comply with the current standards and IETF Best Current Practice (BCP195, RFC5246 at the time of writing) with the additional restrictions enumerated in Authenticate Node [ITI-19] section Mutual TLS / Server-side TLS Option.

For details see RFC7525: https://www.rfc-editor.org/rfc/rfc7525.

Server-side TLS Option

Server-side TLS option is a one-way authentication mechanism where the server identifies itself to the client application via the Transport Layer Security (TLS) protocol with the use of digital certificates. Such authentication mechanism is used in the HTTPS protocol.

The client application uses other means for identification, typically OAuth2/OIDC.

This option is described in the IUA specification for ATNA profile as STX: HTTPS IUA Option.

Note: The abilities and compliance for an actor that supports Server-side TLS option are the same as those described in section Mutual TLS Option.

FQDN Validation of Server Certificate Option

See sections Machine to Machine Authentication and FQDN Validation of Server Certificate.

Note: IETF Best Current Practice BCP195 recommends, but does not require, FQDN validation.

When an actor implements this option, it need not be capable of functioning without this validation.

CA:Sec Required Actor Groupings

There are no required groupings for CA:Sec.

CA:Sec Overview

CA:Sec specifies foundational components that are focused on the security aspect of an overall security system. These are:

- Node Authentication
- Secure Communications

Node authentication enables communications participants to:

- Confirm that the server is indeed the authorized server system.
- Confirm that the client application is indeed an authorized client.

This enables the use of system or machine-level access controls that limit access to only authorized and authenticated machines. The local governance policies will determine whether machine level access control rules are used.

Secure Communications are provided using TLS. TLS provides mutual authentication, reliable message transport and private communication through data encryption. Different forms of encryption can be negotiated to protect the data in transit. CA:Sec permits payload encryption for those sites that wish to implement an additional layer of protection.

CA:Sec does not restrict implementations and deployments to only use the CA:Sec specified methods. For interoperability reasons, TLS must be implemented and available to be configured. The RFC7525 “Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)” covers many configuration options. Deployments often follow these recommendations and make them part of their security policies. A deployment’s security analysis may lead to different choices. Therefore, it is important that implementations allow configuring different protocol versions, algorithms, etc.

Concepts

Cybersecurity activities include a variety of operational, technical, and administrative activities. These are specified in some areas by law or regulation. All of the laws and regulations are consistent in requiring an overall governance model, various technical tools, and operational behaviors. The technical tool requirements always include system authentication, user authentication, event logging (audit), and telecommunications encryption. They also include many other technical features regarding access control, confidentiality, user administration, backups, etc. There are typically also significant operational and administrative requirements.

CA:Sec specifies node authentication and telecommunications encryption. It assumes that the CA:Sec actors will be installed into an environment that complies with all the other governance requirements. Compliance with CA:Sec alone, without also performing the other cybersecurity activities, is not sufficient to provide adequate cybersecurity.

Governance

The specific requirements for cybersecurity vary for different locations and purposes. The overall goals always include protecting confidentiality of data, integrity of data and systems, and availability of data and systems. The requirements affect:

- administrative policies, such as the policies to be used when authenticating and provisioning a new user,
- technical capabilities, such as performing real time access control, and
- operational activities, such as maintaining backup facilities and having continuity of service plans.

It is not practical or reasonable for CA:Sec to profile those requirements. They are too varied and cover much more than just interoperability of systems.

CA:Sec assumes that governance is established that is similar to the recommendations found in the NIST and ISO Frameworks:

- NIST Cybersecurity Framework PR.PT.1, https://www.nist.gov/cyberframework
- NIST SP 800-53 Rev4 AU Family https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final
- ISO 27001:2013 A.12.4.1, A.12.4.2, A.12.4.3, A.12.4.4, A.12.7.1 https://www.iso.org/standard/54534.html

Authentication

User Authentication

The remote Secure Application validates that the user who requests access has been (locally or remotely) authenticated and is authorized to use the service requested.

For locally authenticated users, authentication must take place prior to the creation of a secure tunnel – resources should not be used prior to local user authentication.

CA:Sec does not specify the specific method for user authentication. IHE has profiles that specify particular kinds of user authentication. These can be used, as can other non-IHE methods for user authentication. For details, see Security Considerations.

Machine to Machine Authentication

CA:Sec specifies that connections between machines be authenticated and use TLS. The TLS machine authentication is based upon the use of public and private certificates. This is the method used to authenticate many sensitive transactions on the Internet.

Unlike the typical Internet browser setup, within a healthcare setting:

- Individual direct comparison for validation of certificates can be practical and appropriate. For example, it can be reasonable to use direct comparison and provide the public certificate between two secure applications.
- Chain of trust signed certificates can be practical and appropriate. It can be reasonable to have a hospital security system provide the trusted root authority for authenticating that a particular machine is an authenticated member of the hospital network.
- The commonly used root certificate authorities for browsers are much less likely to be appropriate for a chain of trust method. Their certificate policies are designed for financial risk reduction, not healthcare system authentication.

A means must be provided to install the required certificates to any CA:Sec implementation so that the systems can be configured to match the local governance. The common browser root certificate list is not sufficient. The particular implementation will ultimately make final ruling regarding the detailed requirements for certificates that are needed for a specific system.

CA:Sec Security Considerations

Compliance with CA:Sec alone, without also performing the other cybersecurity activities, is not sufficient to provide adequate cybersecurity.

There are many security-related aspects that play a critical role in providing adequate cybersecurity to computer systems, such as User Authentication, Authorization, Access Control, Privacy/Consent, Logging, Auditing, Governance and more. While these cybersecurity activities are of utmost importance, these are not included in the focus area of CA:Sec Implementation Guidance. Instead, many of these aspects are covered in other IHE profiles that are specialized in those areas.

To achieve a high degree of cybersecurity, the actors from the various IHE profiles can be grouped together, wherever possible. Alternatively, other non-IHE methods can be used.

CA:Sec Cross Profile Considerations

CA:Aud

The Canadian Audit Trail (CA:Aud) Implementation Guidance provides support for Event Logging for Auditing. To allow for these features, CA:Sec actors can be grouped with CA:Aud actors.

If the grouping is in place, an actor from CA:Sec shall implement the required transactions and/or content modules in CA:Sec ***in addition to all*** the transactions required for the grouped actor (Column 2).

CA:Sec Actor	**Actor(s) to be grouped with**
Secure Application	CA:Aud / Audit Creator

CA:Sec actors, when grouped with Audit Creator, shall use the Record Audit Event [ITI-20] transaction to send audit event log messages to an Audit Record Repository.

Alternatively, other non-IHE methods can be used to record audit messages, that do not require grouping with CA:Aud actors.

IUA

The Internet User Authorization (IUA) Profile provides support for user authentication, app authentication, and authorization decisions. To allow for these features, CA:Sec actors can be grouped with IUA actors.

If the grouping is in place, an actor from CA:Sec shall implement the required transactions and/or content modules in CA:Sec ***in addition to all*** the transactions required for the grouped actor (Column 2).

CA:Sec Actor	**Actor(s) to be grouped with**
Local Secure Application (Client)	IUA / Authorization Client
Remote Secure Application (Server)	IUA / Resource Server

This option is described in the IUA specification for ATNA profile as STX: HTTPS IUA Option.

Actors that support this option utilize server-side authenticated TLS (also known as https) to authenticate the server to the client and provide communications integrity and encryption.

This configuration utilizes ATNA Profile server-side TLS (https) to authenticate the server to the client and provide communications integrity and encryption; and the IUA Profile to authenticate the client application to the server (IUA Resource Server).

- TLS shall be server side authenticated, and may be client authenticated
- TLS shall be compliant with BCP195
- Local Secure Node or Secure Application shall reject connections that are not https, and may enforce other policies
- Remote Secure Node or Secure Application shall reject connections that do not carry a valid IUA token, and may enforce other policies

CA:Sec Transactions

The purpose is to define the details and constraints for the transactions directly involved in the CA:Sec Implementation Guidance.

Authenticate Node [ITI-19]

This section corresponds to transaction [19] of the IHE ITI Technical Framework. Transaction [ITI-19] is used by the Secure Application Actor.

Scope

In the Authenticate Node transaction, a Secure Application presents its identity to another Secure Application. This authentication can be one of the following:

- Mutual (two-way) authentication, where both the server and the client identify themselves using the Authenticate Node [ITI-19] transaction. After this mutual authentication, other secure transactions may take place through this secure pipe between the two nodes. In addition, the Secure Application authenticates the identity of the user who requests access to the node. This user authentication is a local operation that does not involve communication with a remote node.
- Server-side (one-way) authentication, where only the server identifies itself using the Authenticate Node [ITI-19] transaction. The client application uses other means for identification, typically OAuth2/OIDC.

Messages

Note: This diagram does not imply sequencing of Authentication Node and Local User Authentication.

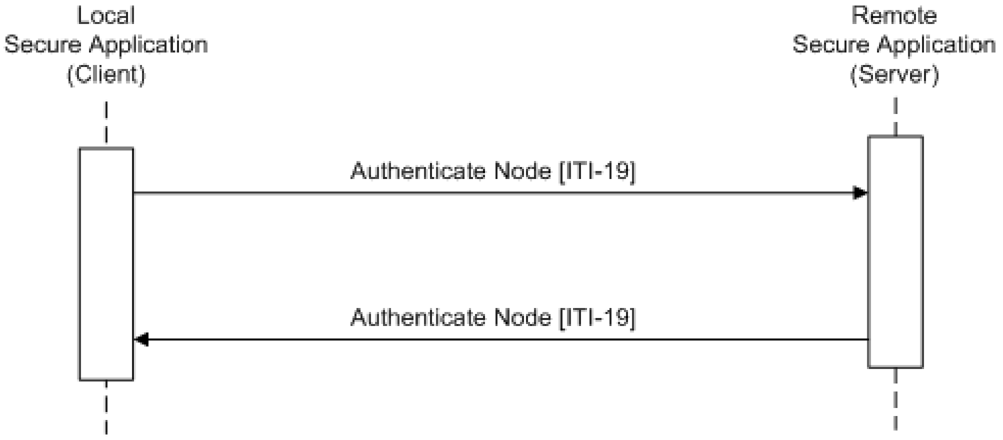

Mutual authentication (both client and server authenticate themselves)

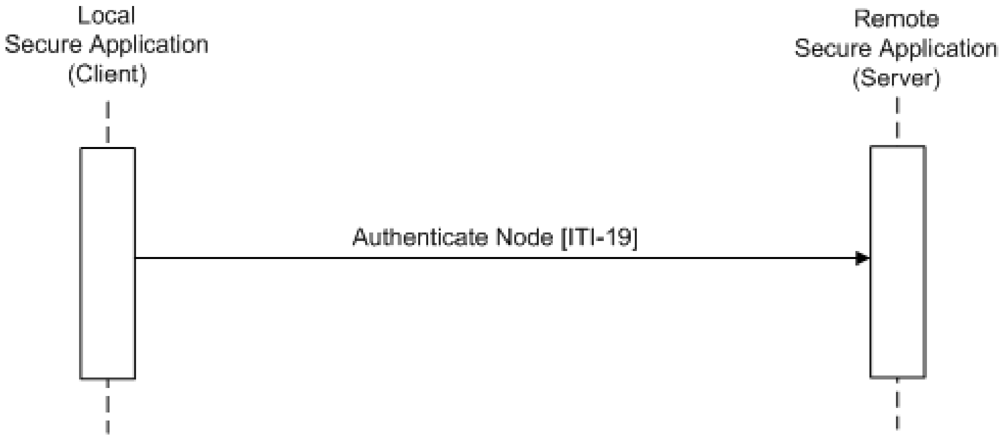

Server-side authentication (only the server authenticates itself)

Trigger Events

- Mutual authentication

The Local Secure Application (Client) initiates the bi-directional authentication process with the Remote Secure Application when information exchange between the two nodes is requested. The first transaction shall be the Authenticate Node transaction, and all subsequent PHI transactions performed by IHE actors shall be secure transactions. This authentication process is needed when a secure connection is established.

The Secure Application shall always apply the Authenticate Node process to every connection.

- Server-side authentication

The Local Secure Application (Client) initiates the server-side authentication process with the Remote Secure Application when information from the server node is requested.

Message Semantics

The Authenticate node transaction involves the exchange of certificates representing the identities of the nodes.

The cipher suites specified here set a baseline to ensure that interoperability is possible. This baseline shall not prohibit more secure configurations from being used. The actual cipher suites will be negotiated in order and according to local policies, with the most secure configurations preferred.

Certificate Validation

The local organization will make the choice of what mixture of chain of trust and direct comparison is used to authenticate and also authorize communications. This may be entirely based on chaining trust to selected Certificate Authorities (CAs), entirely based upon provision of node certificates for direct comparison, or a mixture of both.

Note: The CAs used for CA:Sec chain of trust will be different than the default browser trusted list of CAs used for authenticating internet web servers. A worldwide CA, such as VeriSign, is not generally trusted to determine which individual nodes within an organization should and should not communicate patient identifiable information.

When Authenticating the Remote Secure Application, the Local Secure Application:

- Shall be able to perform certificate validation based on signature by a trusted CA (see section Chain to a trusted certificate authority) and
- Shall be able to perform direct certificate validation to a set of trusted certificates (see section Direct certificate validation)

It may reject communications when the certificate validation fails or may restrict communications to only that which is appropriate for an unidentified other party.

Chain to a trusted certificate authority

The Secure Application:

- Shall provide the means for configuring which CAs are trusted to authenticate node certificates for use in a chain of trust. These CAs shall be identified by means of the public signing certificate for the signing CA.
- Shall support digital certificates encoded using both Deterministic Encoding Rules (DER) and Basic Encoding Rules (BER).
- Shall accept communications for which there is a certificate that is signed by a CA that is listed as a trusted signing authority.

Additional security considerations for API transactions:

- There may be other methods with respect to authorization, e.g., API key or certificate issued by the app.

Direct certificate validation

The Secure Application:

- Shall provide means for installing of the required certificates, for example, via removable media or network interchange (where the set of trusted certificates can be a mixture of CA signed certificates and self-signed certificates).
- Shall support digital certificates encoded using both DER and BER.
- Shall accept communications for which there is a certificate configured as acceptable for direct certificate validation.

Other Certificate requirements

The Secure Application shall not require any specific certificate attribute contents, nor shall it reject certificates that contain unknown attributes or other parameters. Note that for node certificates the Common Name (CN) often is a hostname, attempting to use this hostname provides no additional security and will introduce a new failure mode (e.g., DNS failure).

The certificates used for mutual authentication shall be X.509 certificates based on either:

- RSA key with key length in the range of 2048-4096 bit, where the key length chosen is based on local site policy and as per minimum accepted by today's standards (NIST SP 800-57, FIPS140-3), or
- BCP195 certificate recommendations.

Maximum expiration time acceptable for certificates should be defined in the applicable security policy. The IHE Technical Framework recommends a maximum expiration time of 2 years.

The method used to determine whether a node is authorized to perform transactions is not specified. This may be use of a set of trusted certificates, based on some attribute value contained in the certificates, access control lists, or some other method. Using a certificate chain back to an external trusted certificate authority to determine authorizations is strongly discouraged.

FQDN Validation of Server Certificate Option

The FQDN Validation of Server Certificate Option applies the rules presented in RFC6125 when a client authenticates the server using an X.509 certificate in the context of Transport Layer Security (TLS).

A client, who is validating a server's identity, shall validate that the reference identifier present in a subjectAltName entry of type DNS-ID matches the source domain of the server, per RFC6125 Section 6. Note that the rules described in RFC6125 Section 6 require the validation to be performed based on the input source and the DNS-ID fully qualified domain name.

In an environment where clients have implemented this option, a server's X.509 certificate shall contain a subjectAltName entry of type DNS-ID, per RFC6125 Section 4.

All Connections carrying Protected Information (PI) using TLS

Mutual TLS / Server-side TLS Option

An actor using the TLS floor Option:

- Shall be able to comply with BCP195. This implies that the implementation:
  - Utilizes the framework and negotiation mechanism specified by the Transport Layer Security protocol.
  - Supports current version of TLS or higher
- Shall also be able to restrict to use current version of TLS or higher.
- Shall also support the following cipher suites if using TLS 1.2:
  - TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
  - TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
  - TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
  - TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- Shall also support the following cipher suites if using TLS 1.3:
  - TLS_AES_256_GCM_SHA384
  - TLS_AES_128_GCM_SHA256
- Shall reject negotiation of any cipher suites that have been identified as providing weak security, or less than 128-bit encryption. A weak cipher is defined as an encryption/decryption algorithm that uses a key of insufficient length, typically less than 128 bits. Examples of weak algorithms are: Data Encryption Standard (DES), Electronic Codebook (ECB) and Cipher Block Chaining (CBC).

Additional cipher suites of similar or greater cryptographic strength may be supported.

Referenced Standards

IETF:

- [RFC6125] - Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <https://www.rfc-editor.org/info/rfc6125>.

- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <https://www.rfc-editor.org/info/rfc5246>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <https://www.rfc-editor.org/info/rfc7525>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.
- [BCP195] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, MAY 2015. Moriarty, K. and S. Farrell, "Deprecating TLS 1.0 and TLS 1.1", BCP 195, RFC 8996, March 2021. <https://www.rfc-editor.org/info/bcp195>
- ATNA - Audit Trail and Node Authentication https://wiki.ihe.net/index.php/Audit_Trail_and_Node_Authentication

ITU-T:

- Recommendation X.509 (03/00). “Information technology - Open Systems Interconnection - The directory: Public-key and attribute certificate frameworks"

Audit Considerations

To add support for audit, CA:Sec actors can be grouped with CA:Aud actors (see section Cross Profile Considerations). Alternatively, other non-IHE methods can be used to record audit messages, that do not require grouping with CA:Aud actors.

To enable audit, a Secure Application shall detect and report the events defined in table below. Additional events may be reported.

Audit Event Trigger	**Description**
Actor-start-stop	Startup and shutdown of any actor. Applies to all actors. Is distinct from hardware powerup and shutdown.
Mobile-machine-event	Mobile machine joins or leaves secure domain.
Node-Authentication-failure	A secure application authentication failure has occurred during TLS negotiation, e.g., invalid certificate.

Security Alert	Security Administrative actions create, modify, delete, query, and display the following: Configuration and other changes, e.g., software updates that affect any software that processes protected information. Hardware changes may also be reported in this event. 1. Security attributes and auditable events for the application functions used for patient management, clinical processes, registry of business objects and methods (e.g., WSDL, UDDI), program creation and maintenance, etc. 2. Security domains according to various organizational categories such as entity-wide, institutional, departmental, etc. 3. Security categories or groupings for functions and data such as patient management, nursing, clinical, etc. 4. The allowable access permissions associated with functions and data, such as create, read, update, delete, and execution of specific functional units or object access or manipulation methods. 5. Security roles according to various task-grouping categories such as security administration, admissions desk, nurses, physicians, clinical specialists, etc. It also includes the association of permissions with roles for role-based access control. 6. User accounts. This includes assigning or changing password or other authentication data. It also includes the association of roles with users for role-based access control, or permissions with users for user-based access control. 7. Unauthorized user attempt to use security administration functions. 8. Audit enabling and disabling. 9. User authentication revocation. 10. Emergency Mode Access (aka Break-Glass) Security administration events should always be audited.
User Authentication	This message describes the event of a user log on or log off, whether successful or not. No Participant Objects are needed for this message.

Canadian Audit Trail (CA:Aud) Implementation Guidance

The **CA:Aud (Canadian Audit Trail)** Implementation Guidance specifies the foundational elements needed to perform event logging for auditing purposes.

CA:Aud is based on the ATNA profile and aims to bring improvements via loose coupling, and high cohesion, with focus on auditing using modern formats and technologies.

CA:Aud defines capabilities to record, store and retrieve audit messages in FHIR format using RESTful operations.

For more details see IHE documentation ATNA profile, RESTful ATNA Supplement, ATNA ITI-20 and ATNA ITI-81 transaction documentation.

Note: CA:Aud is meant to be used in a secure environment, that is compliant with critical security and privacy requirements that together provide adequate cybersecurity to the overall system, such as Secure Communication, User Authentication, Authorization, Access Control, Privacy/Consent, Governance and more. Audit logs from other platforms such as web server, operating system and database are also expected.

CA:Aud Description

The purpose is to define the actors directly involved in the CA:Aud Implementation Guidance, the transactions between them, and the options and groupings of the actors with other IHE profiles.

CA:Aud Actors/Transactions

The diagram below shows the actors directly involved in CA:Aud and the relevant transactions between them.

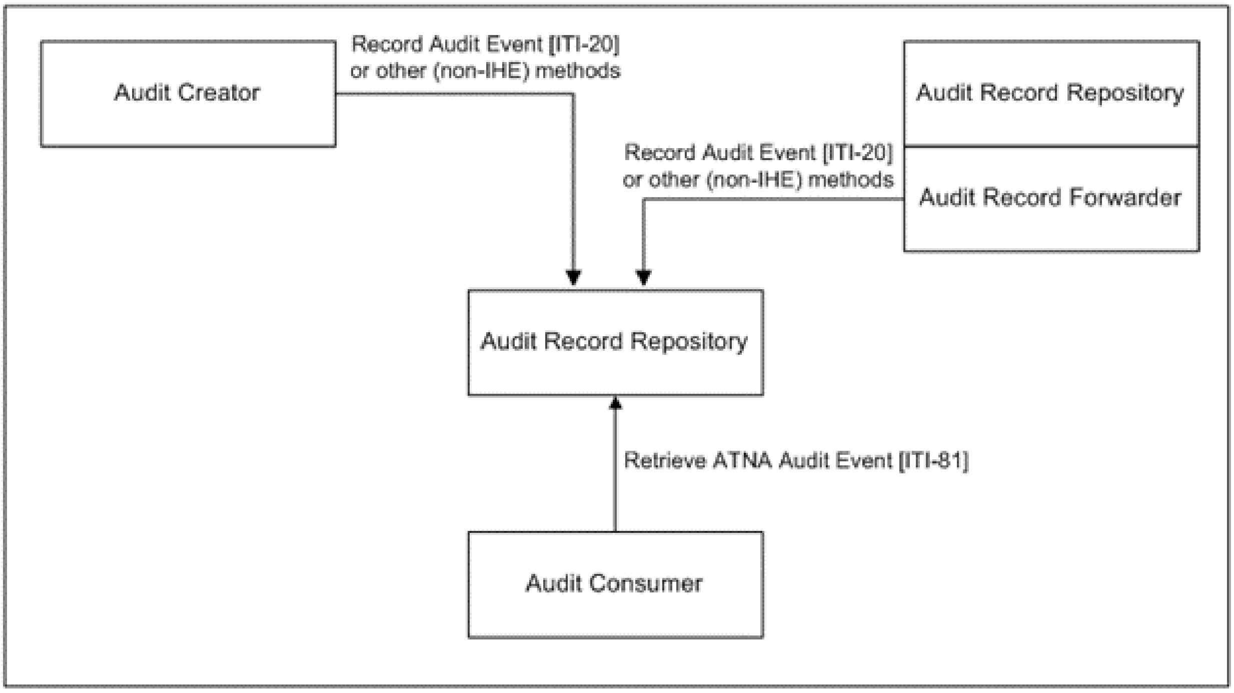
Audit Creator
Record Audit Event [ITI-20] or other (non-IHE) methods
Audit Record Repository
Audit Record Forwarder
Record Audit Event [ITI-20] or other (non-IHE) methods
Audit Record Repository
Retrieve ATNA Audit Event [ITI-81]
Audit Consumer

The table below lists the transactions for each actor directly involved in CA:Aud.

To claim compliance with this CA:Aud, an actor shall support all required transactions (labeled “R”) and may support the optional transactions (labeled “O”).

Actors	**Transactions**	**Optionality**
Audit Creator	Record Audit Event [ITI-20]	O (Note 1)
Audit Record Repository	Record Audit Event [ITI-20]	O (Note 1)
	Retrieve ATNA Audit Event [ITI-81]	O (Note 2)

Audit Record Forwarder	Record Audit Event [ITI-20]	O (Note 1)
Audit Consumer	Retrieve ATNA Audit Event [ITI-81]	R

Note 1: The audit events must be recorded using the IHE Record Audit Event [ITI-20] with FHIR Feed option or other (IHE or non-IHE) methods.

Note 2: This transaction is required if the Audit Record Repository is central.

Note that if other (IHE or non-IHE) methods are used to record the audit events, the messages must be converted into the FHIR format that is expected by the Audit Consumer for Retrieve ATNA Audit Event [ITI-81] transaction.

Actor Descriptions and Requirements

Audit Creator

The Audit Creator creates AuditEvent records in FHIR format as specified in Record Audit Event [ITI-20] and sends these records to the Audit Record Repository.

The Audit Creator shall:

1. Perform only secure transactions to or from the node.
2. Provide sufficient authentication methods, based on risk assessment, to ensure that only authorized users access the Audit Creator.
3. Detect and report a Record Audit Event as specified in Record Audit Event [ITI-20] or by other (IHE or non-IHE) means for:
    - all the activity-related events for the Audit Creator
    - all transaction-related events for the Audit Creator

Audit Record Repository

The Audit Record Repository receives event audit reports and stores them. It may be part of a federated network of repositories. It is expected to have analysis and reporting capabilities, but those capabilities are not specified as part of CA:Aud. CA:Aud also does not specify the capacity of an Audit Record Repository, because the variety of deployment needs makes it impractical to set requirements for the event report volume or capacity needed.

The Audit Record Repository shall support:

1. Recording and storing of Record Audit Events via Record Audit Event [ITI-20] or by other (IHE or non-IHE) means.
2. Search capabilities as defined in Retrieve ATNA Audit Event [ITI-81].
3. Local security and privacy service protections and user access controls.

The Audit Repository may support:

1. The FHIR Feed audit transport mechanism specified in Record Audit Event [ITI-20].
2. Receipt of an IHE-specified audit message format. Note that the message format is extensible to include both future IHE specifications (e.g., audit requirements for new IHE transactions) and private extensions.
3. Other (IHE or non-IHE) transport mechanisms and message formats for audit records.

Audit Record Forwarder

The Audit Record Forwarder is grouped with an Audit Record Repository, and forwards selected audit messages that are received by the Audit Record Repository. It may filter these messages and forward them selectively. It may forward to multiple different Audit Record Repositories.

The Audit Record Forwarder shall:

1. Be grouped with an Audit Record Repository.

2. Filter and forward messages as they arrive.
3. Be configurable to forward messages to destination Audit Record Repositories.

Audit Consumer

The Audit Consumer queries an Audit Record Repository for CA:Aud audit records using CA:Aud audit record content. Subsequent processing of the query result is not defined in this CA:Aud.

CA:Aud Actor Options

For each actor in CA:Aud, the options (labeled "R") shall be selected, and the options (labeled "O") may be selected.

Actor	**Options**	**Optionality**
Audit Creator	FHIR Feed	O (Note 1)
Audit Record Repository	FHIR Feed	O (Note 1)
	Retrieve Audit Message	O (Note 2)
Audit Record Forwarder	FHIR Feed	O (Note 1)
Audit Consumer	Retrieve Audit Message	R

Note 1: The recording of audit events is mandatory, however CA:Aud does not enforce the means that are used to record the audit events into the Audit Record Repository. Audit records must be recorded using the IHE Record Audit Event [ITI-20] with FHIR Feed option or by other (IHE or non-IHE) methods. The FHIR Feed option is recommended.

Note 2: This transaction is required if the Audit Record Repository is central.

Note that if other (IHE or non-IHE) methods are used to record the audit events, the messages must be converted into the FHIR format that is expected by the Audit Consumer for Retrieve ATNA Audit Event [ITI-81] transaction.

FHIR Feed Option

The audit message transport happens via FHIR Feed, that enables sending CA:Aud audit records using RESTful capabilities and FHIR resources.

The Audit Record Repository shall implement two RESTful interactions, Send Audit Resource and Send Audit Bundle, as defined in the Record Audit Event [ITI-20] transaction.

An Audit Creator or Audit Record Forwarder shall at least support one of the two RESTful interactions, Send Audit Resource and Send Audit Bundle, as defined in the Record Audit Event [ITI-20] transaction.

FHIR Feed is the recommended option, as it provides the appropriate FHIR format that can be consumed by the Audit Consumer actor using Retrieve ATNA Audit Event [ITI-81] transaction.

Retrieve Audit Message Option

The Retrieve Audit Message Option enables search requests for audit records based upon message contents.

An Audit Record Repository that supports this option shall implement the Retrieve ATNA Audit Event [ITI-81] transaction.

The [ITI-81] transaction is a RESTful search from an Audit Consumer to an Audit Record Repository (ARR) using FHIR resources. The search response will reflect the contents of the data storage at the time of the search. CA:Aud does not specify the criteria for message selection, archival, retention interval, etc. These are set by local policy and often vary for different Audit Record Repositories.

CA:Aud Required Actor Groupings

An actor from CA:Aud shall implement the required transactions and/or content modules in CA:Aud ***in addition to all*** the transactions required for the grouped actor (Column 2).

CA:Aud Actor	**Actor(s) to be grouped with**
Audit Record Forwarder	CA:Aud / Audit Record Repository

CA:Aud Overview

CA:Aud specifies foundational components that are focused on:

- Event Logging (Audit)

Successful implementation of CA:Aud also requires the existence and support of:

- Secure Communication
- System Security Services
- Access control
- Privacy and Security Governance

For event audit logging, CA:Aud specifies:

- A standard schema for encoding a reported event
- Standard events to be reported:
  - Events that are related to system activities, e.g., “Login Failure”.
  - Events that are related to IHE transactions. These are described in the technical framework sections that describe the transaction.
- Event reporting messages in FHIR format using RESTful operations.
- An Audit Record Repository for collecting and reporting on the event audit logs.

Concepts

CA:Aud assumes that the actors will be installed into an environment that complies with all the security, privacy, and governance requirements.

Governance

The specific requirements for cybersecurity vary for different locations and purposes. The overall goals always include protecting confidentiality of data, integrity of data and systems, and availability of systems.

It is not practical or reasonable for CA:Aud to profile those requirements. They are too varied and cover much more than just interoperability of systems.

Event Logging

CA:Aud event audit logging is intended to provide a surveillance logging function. This means that it captures:

- All security events that are detected.
- A full set of activity and transaction events describing ongoing operations. These are used to establish a baseline for what is normal operation and are monitored for deviations from that baseline. The level of

detail is subject to judgment. Details that do not matter in terms of establishing what is normal are left out, especially if they would reveal PHI.

The event logging is not designed for:

- Detailed forensic analysis, such as will be performed when surveillance reveals suspicious activity or after a security event is detected. This often needs to be at a level of detail that involves specific design aspects of specific products. CA:Aud expects that there is a forensic level log for products, and that those products document the design and specific details of their event reports. The forensic log may also use the CA:Aud schema and transactions, or it may be different.
- Workflow performance analysis log, such as is typical in tightly coordinated system controls. The CA:Aud events were chosen for privacy and security surveillance, not for system or staff performance purposes. A workflow analysis log may also use the CA:Aud schema and transactions, or it may be different.

Events

Activity

CA:Aud defines events related to activities of the IHE actors and system components that are grouped with a secure actor. These include events such as system startup, user login (both success and failure), access control violation, etc. CA:Aud requires that these be detected and reported.

These events are described in Record Audit Event [ITI-20], see sections Send Audit Resource Request Message, Trigger Events. Additional reportable events are often identified for specific events in other IHE profiles and are documented in those profiles or transactions, or they may be specified by local law, regulation, or policy.

Transaction

IHE profiles that define transactions may define events and specify the event reporting structure for those events.

CA:Aud Security Considerations

There are many security-related aspects that play a critical role in providing adequate cybersecurity to computer systems, such as User Authentication, Authorization, Access Control, Privacy/Consent, Logging, Auditing, Governance and more. While these cybersecurity activities are of utmost importance, these are not included in the focus area of CA:Aud Implementation Guidance. Instead, many of these aspects are covered in other IHE profiles that are specialized in those areas.

To achieve a high degree of cybersecurity, the actors from the various IHE profiles can be grouped together, wherever possible. Alternatively, other (IHE or non-IHE) methods can be used.

Some basic concepts are described in CA:Aud Overview.

CA:Aud defines transactions for the Audit Record Repository that enable sharing of sensitive information related to patients and systems.

In many implementations and projects, Audit Record Repository have been considered a “black-box” able to store relevant information for security and monitoring purposes. Those systems have not historically been designed to provide external access to stored records.

Security Officers and System Architects should consider this and analyze the risks of disclosing data stored in the Audit Record Repository. The Retrieve ATNA Audit Event [ITI-81] transaction define how to search the audit records stored in FHIR format captured using Record Audit Event [ITI-20] transaction.

Accordingly, access control mechanisms on the CA:Aud actors and queries are strongly recommended. The Internet User Authorization (IUA) Profile should be considered for the authorization controls. The CA:Aud Audit Record Repository can be grouped with an IUA Resource Server to enforce policies and authorization decisions. The Audit Consumer can be grouped with an IUA Authorization Client to provide authorization information to the CA:Aud Audit Record Repository. Access controls should appropriately restrict access to audit records.

The Retrieve CA:Aud Audit Event [ITI-81] transaction may involve the disclosure of sensitive information. Logging this retrieval transaction as a query event is appropriate (see section Retrieve ATNA Audit Event [ITI-81] Security Audit Considerations).

Additional Security Considerations are described in Z.8 Mobile Security Considerations of RESTful ATNA Supplement.

CA:Aud Cross Profile Considerations

CA:Sec

The Canadian Network Security (CA:Sec) Implementation Guidance provides support for secure communication. To allow for these features, CA:Aud actors can be grouped with CA:Sec actors.

If the grouping is in place, an actor from CA:Aud shall implement the required transactions and/or content modules in CA:Aud ***in addition to all*** the transactions required for the grouped actor (Column 2).

CA:Aud Actor	**Actor(s) to be grouped with**
Audit Creator	CA:Sec / Secure Application
Audit Record Repository	CA:Sec / Secure Application
Audit Record Forwarder	CA:Sec / Secure Application

CA:Aud actors, when grouped with CA:Sec Secure Application, shall use the Authenticate Node [ITI-19] transaction to ensure secure communication between actors.

CT

The Consistent Time (CT) Profile provides a means to ensure that the system clocks and time stamps of the computers in a network are synchronized. To allow for these features, CA:Aud actors can be grouped with CT actors.

If the grouping is in place, an actor from CA:Aud shall implement the required transactions and/or content modules in CA:Aud ***in addition to all*** the transactions required for the grouped actor (Column 2).

CA:Aud Actor	**Actor(s) to be grouped with**
Audit Creator	CT / Time Client
Audit Record Repository	CT / Time Client
Audit Record Forwarder	CT / Time Client

CA:Aud Transactions

The purpose is to define the details and constraints for the transactions directly involved in the CA:Aud Implementation Guidance.

Record Audit Event [ITI-20]

This section corresponds to the Record Audit Event [ITI-20] transaction of the IHE IT Infrastructure Technical Framework. This transaction is used to report auditable events to an Audit Record Repository.

Scope

This transaction is used to report auditable events to an Audit Record Repository using FHIR Feed option.

Actor Roles

Actor	**Role**
Any actor grouped with the Audit Creator	Create an audit record and transmit this record to the Audit Record Repository.
Audit Record Repository	Receive an audit record from the Audit Record Creator and store this for audit purposes.
Audit Record Forwarder	Forward an audit record to Audit Record Repositories.

Referenced Standards

HL7 FHIR	Release 4 http://hl7.org/fhir/R4/index.html
RFC4627	The application/json Media Type for JavaScript Object Notation (JSON)

Messages

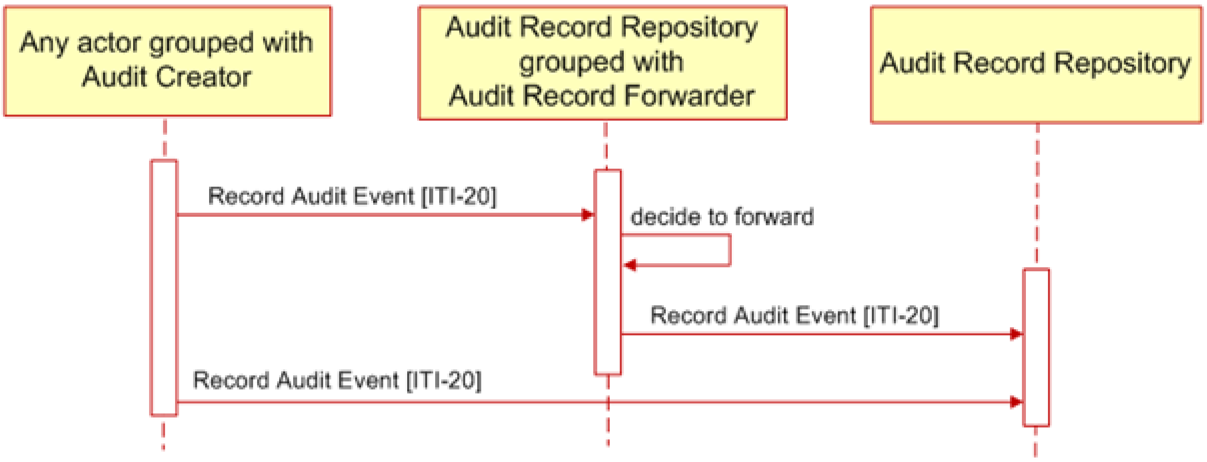
Any actor grouped with Audit Creator
Audit Record Repository grouped with Audit Record Forwarder
Audit Record Repository
Record Audit Event [ITI-20]
decide to forward
Record Audit Event [ITI-20]
Record Audit Event [ITI-20]

Note 1: Any actor initiating [ITI-20] may send to more than one Audit Record Repository.

Note 2: The Audit Repository that receives an [ITI-20] transaction may or may not be grouped with an Audit Record Forwarder. This diagram does not show a chain of forwarding between actors.

In the context of CA:Aud, Transaction [ITI-20] defines two different interactions that can be used for auditing:

1. The “Send Audit Resource Request Message - FHIR Feed Interaction” is used for auditing a single FHIR AuditEvent Resource using RESTful protocol (see section Send Audit Resource Request Message).
2. The “Send Audit Bundle Request Message - FHIR Feed Interaction” is used for auditing a bundle of FHIR AuditEvent Resources using RESTful protocol (see section Send Audit Bundle Request Message).

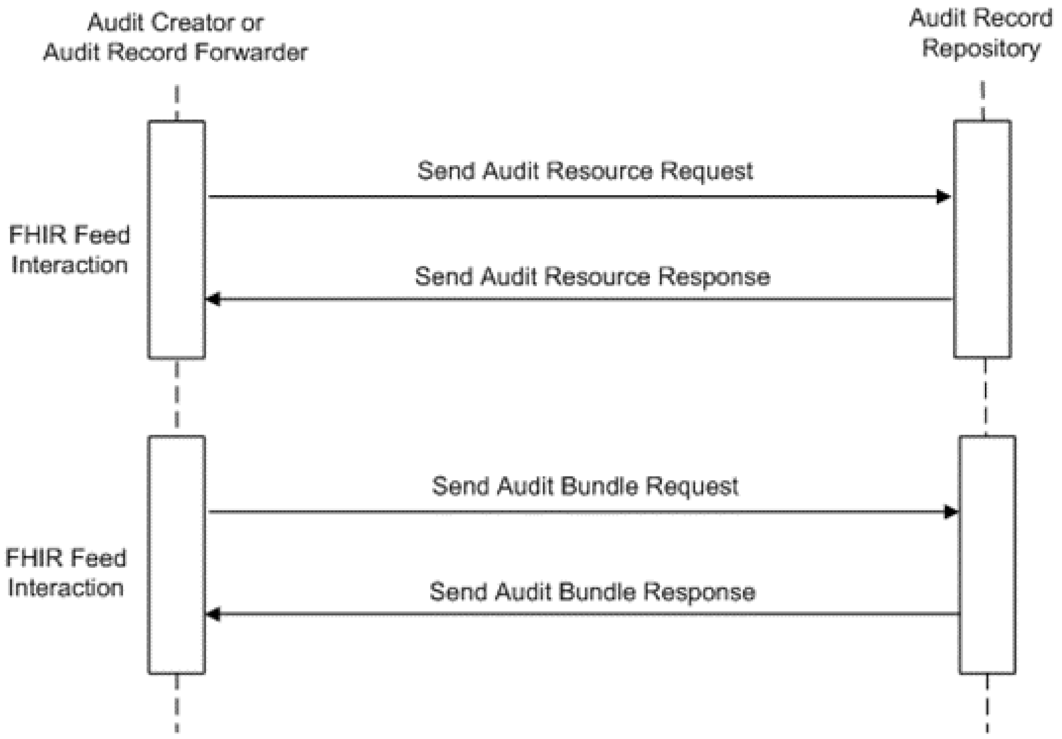
Audit Creator or
Audit Record Forwarder
Audit Record
Repository
FHIR Feed
Interaction
Send Audit Resource Request
Send Audit Resource Response
FHIR Feed
Interaction
Send Audit Bundle Request
Send Audit Bundle Response

Send Audit Resource Request Message – FHIR Feed Interaction

An actor that is grouped with an Audit Creator or Audit Record Forwarder, detects an event that should be reported and uses the Send Audit Resource Request message to send a report about the event to an Audit Record Repository.

An Audit Creator or Audit Record Forwarder, that supports the FHIR Feed Interaction, uses this message to post a single AuditEvent Resource to the Audit Record Repository using a FHIR create interaction (see https://www.hl7.org/fhir/R4/http.html#create).

Trigger Events

This message is sent when an actor that is grouped with an Audit Creator or Audit Record Forwarder needs to post a single AuditEvent Resource to the Audit Record Repository.

There are two trigger events:

1. An Audit Creator detects an event that should be reported to the Audit Record Repository. This transaction does not specify all the policies or reasons for reporting events. They may be specified in other IHE profiles, they may be specified by local law or regulation, or they may be specified by local policy.
2. An Audit Record Forwarder determines that a received AuditEvent Resource should be sent to another Audit Record Repository. This transaction does not specify what rules or policies determine whether an AuditEvent Resource should be forwarded.

An actor in any IHE profile, when grouped with an Audit Creator, shall be able to report the events defined in table below. Additional reportable events are often identified for specific events in other IHE profiles and are documented in that profile or transaction.

Audit Record trigger events:

Audit Event Trigger	**Description**
Actor-start-stop	Startup and shutdown of any actor. Applies to all actors. Is distinct from hardware powerup and shutdown.
Audit-Log-Used	The audit trail repository has been accessed or modified by something other than the arrival of audit trail messages.
Begin-storing-instances	Begin storing SOP Instances for a study. This may be a mix of instances.
Instances-deleted	SOP Instances are deleted from a specific study. One event covers all instances deleted for the particular study.
Instances-Stored	Instances for a particular study have been stored on this system. One event covers all instances stored for the particular study.
Mobile-machine-event	Mobile machine joins or leaves secure domain.
Node-Authentication-failure	A secure node authentication failure has occurred during TLS negotiation, e.g., invalid certificate.
Order-record-event	Order record created, accessed, modified or deleted. Involved actors: Order Placer. This includes initial order, updates or amendments, delivery, completion, and cancellation. See note below.
Patient-record-event	Patient record created, modified, or accessed.
PHI-export	Any export of PHI on media, either removable physical media such as CD-ROM or electronic transfer of files such as email. Any printing activity, paper or film, local or remote, which prints PHI.
PHI-import	Any import of PHI on media, either removable physical media such as CD-ROM or electronic transfers of files such as email.
Procedure-record-event	Procedure record created, modified, accessed or deleted.

Query Information	A query has been received, either as part of an IHE transaction, or as part other products functions. For example: 1. Modality Worklist Query 2. Instance or Image Availability Query 3. PIX, PDQ, or XDS Query Notes: The general guidance is to log the query event with the query parameters and not the result of the query. The result of a query may be very large and is likely to be of limited value vs. the overhead. The query parameters can be used effectively to detect bad behavior and the expectation is that given the query parameters the result could be regenerated if necessary.

Security Alert	Security Administrative actions create, modify, delete, query, and display the following: Configuration and other changes, e.g., software updates that affect any software that processes protected information. Hardware changes may also be reported in this event. 1. Security attributes and auditable events for the application functions used for patient management, clinical processes, registry of business objects and methods (e.g., WSDL, UDDI), program creation and maintenance, etc. 2. Security domains according to various organizational categories such as entity-wide, institutional, departmental, etc. 3. Security categories or groupings for functions and data such as patient management, nursing, clinical, etc. 4. The allowable access permissions associated with functions and data, such as create, read, update, delete, and execution of specific functional units or object access or manipulation methods. 5. Security roles according to various task-grouping categories such as security administration, admissions desk, nurses, physicians, clinical specialists, etc. It also includes the association of permissions with roles for role-based access control. 6. User accounts. This includes assigning or changing password or other authentication data. It also includes the association of roles with users for role-based access control, or permissions with users for user-based access control. 7. Unauthorized user attempt to use security administration functions. 8. Audit enabling and disabling. 9. User authentication revocation. 10. Emergency Mode Access (aka Break-Glass) Security administration events should always be audited.
User Authentication	This message describes the event of a user log on or log off, whether successful or not. No Participant Objects are needed for this message.
Study-Object-Event	Study is created, modified, accessed, or deleted. This reports on addition of new instances to existing studies as well as creation of new studies.
Study-used	SOP Instances from a specific study are created, modified or accessed. One event covers all instances used for the particular study.

Message Semantics

An Audit Creator or Audit Record Forwarder shall issue an HTTP request according to requirements defined in the FHIR specification for “create” interaction (http://hl7.org/fhir/R4/http.html#create). The message uses an HTTP POST method to send a FHIR AuditEvent Resource.

The Audit Creator or Audit Record Forwarder shall submit the FHIR AuditEvent Resource in either XML format or JSON format. Values for mime-type of the request message are defined in the ITI TF-2: Appendix Z.6.

An AuditEvent Resource that reflects Audit Message definition defined in IHE Technical Framework shall conform to the requirements below.

FHIR AuditEvent Resource Attribute	**Description**
type	Identifier for a family of the event. For example, a menu item, program, rule, policy, function code, application name or URL. It identifies the performed function.
subtype	Identifier for the category of event.
action	Indicator for type of action performed during the event that generated the audit.
recorded	The time when the event was recorded.
outcome	Indicates whether the event succeeded or failed.
outcomeDesc	A free text description of the outcome of the event.
purposeOfEvent	The purposeOfUse (reason) that was used during the event being recorded.
agent	An actor taking an active role in the event or activity that is logged.
agent.type	Specification of the participation type the user plays when performing the event.
agent.role	The security role that the user was acting under, that come from local codes defined by the access control security system (e.g. RBAC, ABAC) used in the local context.
agent.who	Reference to who this agent is that was involved in the event.
agent.altId	Alternative agent Identifier. For a human, this should be a user identifier text string from authentication system. This identifier would be one known to a common authentication system (e.g. single sign-on), if available.

agent.name	Human-meaningful name for the agent.
agent.requestor	Indicator that the user is or is not the requestor, or initiator, for the event being audited.
agent.policy	The policy or plan that authorized the activity being recorded. Typically, a single activity may have multiple applicable policies, such as patient consent, guarantor funding, etc. The policy would also indicate the security token used.
agent.media	Type of media involved. Used when the event is about exporting/importing onto media.
agent.network.address	An identifier for the network access point of the user device for the audit event.
agent.network.type	An identifier for the type of network access point that originated the audit event.
source	The system that is reporting the event.
source.site	Logical source location within the healthcare enterprise network. For example, a hospital or other provider location within a multi-entity provider group.
source.observer	Identifier of the source where the event was detected.
source.type	Code specifying the type of source where event originated.
entity	Specific instances of data or objects that have been accessed.
entity.what	Identifies a specific instance of the entity. The reference should be version specific.
entity.type	The type of the object that was involved in this audit event.
entity.role	Code representing the role the entity played in the event being audited.
entity.lifecycle	Identifier for the data life-cycle stage for the entity.
entity.securityLabel	Security labels for the identified entity.

entity.name	A name of the entity in the audit event.
entity.query	The query parameters for a query-type entities.
entity.detail	Tagged value pairs for conveying additional information about the entity.
entity.detail.type	The type of extra detail provided in the value.
entity.detail.ValueBase64Binary	The value of the extra Base64Binary detail.

Expected Actions

The Audit Record Repository shall support all the mime-types defined in ITI TF-2: Appendix Z.6.

On receipt of the Send Audit Resource Request message, the Audit Record Repository shall validate the Resources and respond with one of the HTTP codes defined in section Message Semantics.

For the Resource received, the Audit Record Repository may:

- discard the Resource as irrelevant
- retain the Resource in an internal data store
- perform other processing on the Resource

The Audit Record Repository may apply a variety of data retention rules to the data store. This transaction does not specify data retention rules. These usually depend upon the purposes assigned to a specific Audit Record Repository.

The Audit Record Repository shall store any resources that were not discarded and make them available for further search via the Retrieve ATNA Audit Event [ITI-81] transaction.

When the Audit Record Repository is grouped with an Audit Record Forwarder, the Audit Record Forwarder shall:

- apply filtering rules to all AuditEvent Resources received by the Audit Record Repository, and
- forward all AuditEvent Resources that match filters to their configured destinations.

Send Audit Resource Response

The Audit Record Repository responds to the Audit Creator or Audit Record Forwarder using a Send Audit Resource Response message in order to inform the client about the result of the operation.

Trigger Events

When the Audit Record Repository has finished storing the AuditEvent Resource received, it sends this message back to the client acknowledging the result of the request.

Message Semantics

The Audit Record Repository returns an HTTP Status code appropriate to the processing, conforming to specification requirements as specified in https://www.hl7.org/fhir/R4/http.html#create.

If the outcome is a success, the http status code of the response shall be a 2xx code. If the outcome is a failure, the Audit Record Repository shall be capable of returning status codes according to what is defined in https://www.hl7.org/fhir/R4/http.html#create.

The Audit Record Repository can return other status codes 4xx or 5xx in accordance to internal business rules that are out of scope for this transaction.

The Audit Record Repository should be able to handle errors in such a way that audit records intended to be recorded are not lost (e.g., errors due to validation of the message format).

Expected Actions

The Audit Record Repository could return failures. For this reason, it is up to the client to decide what to do with failures that have been returned by the Audit Record Repository.

Send Audit Bundle Request Message – FHIR Feed Interaction

An Audit Creator or Audit Record Forwarder that supports the ATX: FHIR Feed Option uses this message to post a Bundle of AuditEvent Resources to the Audit Record Repository using a FHIR batch interaction (see https://www.hl7.org/fhir/R4/http.html#transaction).

Trigger Events

This message is sent when an Audit Record Forwarder or an actor that is grouped with Audit Creator needs to send multiple events that has been audited to the Audit Record Repository.

There are two trigger events:

1. An Audit Creator detects at least one event that should be reported to the Audit Record Repository. This transaction does not specify all of the policies or reasons for reporting events. They may be specified in other IHE profiles, they may be specified by local law or regulation, or they may be specified by local policy.
2. An Audit Record Forwarder determines that at least one received AuditEvent Resource should be sent to another Audit Record Repository. This transaction does not specify what rules or policies determine whether an AuditEvent Resource should be forwarded.

An actor in any IHE profile, when grouped with an Audit Creator, shall be able to report the events defined in the table in section Send Audit Resource Request Message, Trigger Events. Additional reportable events are often identified for specific events in other IHE profiles and are documented in that profile or transaction.

Message Semantics

An Audit Record Forwarder or an actor that is grouped with Audit Creator shall issue an HTTP request according to requirements defined in the FHIR specification for “batch” interaction (see https://www.hl7.org/fhir/R4/http.html#transaction).

The Audit Record Repository and the client shall both support the “batch” interaction. The message uses an HTTP POST method to submit a FHIR Bundle Resource. The client shall post FHIR resources in either XML format or JSON format. Values for mime-type of the request message are defined in ITI TF-2: Appendix Z.6.

The FHIR Bundle Resource shall contain at least one FHIR AuditEvent Resource (https://www.hl7.org/fhir/R4/auditevent.html).

The element Bundle.entry.request.method shall be POST.

AuditEvent Resources included in the Bundle that reflect Audit Message definitions defined in IHE Technical Framework shall conform to the requirements defined in section Send Audit Resource Request Message Semantics.

Bundle Resource Constraints:

Element & Cardinality	**Constraints**

type [1..1]	Shall be: batch
entry [1..*]	Shall contain at least one AuditEvent Resource
entry.request.method	Shall be: POST

Expected Actions

The Audit Record Repository shall support all the mime-types defined in ITI TF-2: Appendix Z.6.

On receipt of the Send Audit Bundle Resource Request, the Audit Record Repository shall validate Resources included in it and respond with one of the HTTP codes defined in section Send Audit Bundle Response Message Semantics.

For each Resource received in the Bundle, the Audit Record Repository may:

- Discard the Resource as irrelevant.
- Retain the Resource in an internal data store.
- Perform other processing on the Resource.

The Audit Record Repository may apply a variety of data retention rules to the data store. This transaction does not specify data retention rules. These are usually dependent upon the purposes assigned to a specific Audit Record Repository.

The Audit Record Repository shall store any resources that were not discarded and make them available for further search via the Retrieve ATNA Audit Event [ITI-81] transaction.

When the Audit Record Repository is grouped with an Audit Record Forwarder, the Audit Record Forwarder shall:

- Apply filtering rules to all AuditEvent Resources received by the Audit Record Repository, and
- Forward all AuditEvent Resources that match filters to their configured destinations.

Send Audit Bundle Response

The Audit Record Repository sends a Send Audit Bundle Response message in response to a Send Audit Bundle Request.

Trigger Events

When the Audit Record Repository has finished storing the AuditEvent Resources received in the Bundle Resource, it sends back this message to the client acknowledging the result of the request.

Message Semantics

The Audit Record Repository returns an HTTP Status code appropriate to the processing, conforming to specification requirements as specified in https://www.hl7.org/fhir/R4/http.html#transaction-response.

When the Audit Record Repository has processed the request shall return an HTTP response with an overall status code.

To allow the client to know the outcome of the transaction, and the identities assigned to the resources by the Audit Record Repository, the Audit Record Repository shall return a Bundle, with type set to batch-response. Each entry element shall contain a response element with an HTTP Status Code which details the outcome of processing of the request entry.

If no “Prefer” header is specified in the request the server should respond as if it is set to return=minimal; see https://www.hl7.org/fhir/R4/http.html#ops.

If the outcome of the entry is a success, the http status code of the response shall be a 2xx code.

If the outcome of the entry is a failure, the Audit Record Repository shall be capable of returning status codes according to what is defined in https://www.hl7.org/fhir/R4/http.html#create.

The Audit Record Repository can return other status codes 4xx or 5xx in accordance to internal business rules that are out of scope for this transaction.

The Audit Record Repository should be able to handle errors in such a way that audit records intended to be recorded are not lost (e.g., errors due to validation of the message format).

Expected Actions

The Audit Record Repository could return a partial success for the Bundle where some resources succeeded and other not. For this reason, it is up to the client to decide what to do with failures that have been returned by the Audit Record Repository.

Security Considerations

The use of the TLS or HTTPS transport mechanism is recommended because the audit event messages often contain PHI or other sensitive information.

The use of the TLS transport mechanism is not always required because there are other means of protection that may be more appropriate in some situations. The decision to use the UDP transport mechanism should be based upon a security and privacy risk analysis.

The data store within the Audit Record Repository may contain sensitive information, and the Audit Record Repository analysis facilities may allow sensitive queries. It will be a high value target for malicious actors and should be protected accordingly.

The Audit Record Repository is required to generate audit event messages for various kinds of use of the data store and configuration changes. This is specified in section Send Audit Resource Request Message, Trigger Events.

If the AuditEvent Message Option is supported on the Audit Record Repository, update, delete and patch interaction of AuditEvent Resources should be managed by local policies.

Retrieve ATNA Audit Event [ITI-81]

This transaction supports the retrieval of CA:Aud audit record from the Audit Record Repository in accordance with a set of search parameters that determine the retrieved event reports.

This transaction enables an Audit Consumer to search audit events that an Audit Record Repository created via the Record Audit Event [ITI-20] FHIR Feed transaction. If the Audit Record Repository stores audit messages in other formats, it should map the audit messages into the FHIR format that can be consumed by the Retrieve ATNA Audit Event [ITI-81] transaction.

This transaction is a profiling of a standard FHIR search of the AuditEvent Resource.

Scope

The Retrieve ATNA Audit Event transaction is used to search CA:Aud events recorded in an CA:Aud Audit Record Repository. The result of this retrieval is a FHIR bundle of AuditEvent Resources that match with a set of search parameters.

Actor Roles

Actor	Role
Audit Record Repository	Provides storage for CA:Aud audit events and responds to queries for a portion of the stored records.
Audit Consumer	Queries for CA:Aud audit records.

Referenced Standards

RFC2616	IETF Hypertext Transfer Protocol – HTTP/1.1
RFC4627	The application/json Media Type for JavaScript Object Notation (JSON)
RFC6585	IETF Additional HTTP Status Codes
RFC3339	Date and Time on the Internet: Timestamps
HL7 FHIR	Release 4 http://hl7.org/fhir/R4/index.html

Messages

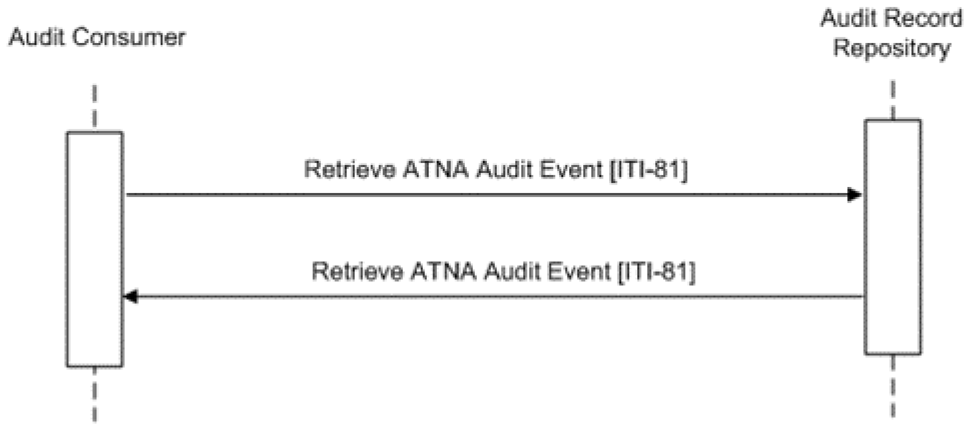
Audit Consumer
Audit Record Repository
Retrieve ATNA Audit Event [ITI-81]
Retrieve ATNA Audit Event [ITI-81]

Retrieve ATNA Audit Events Request Message

This is an HTTP GET parameterized search from an Audit Consumer to an Audit Record Repository. The Audit Record Repository has stored CA:Aud audit records received via Record Audit Event [ITI-20] transactions. Those messages, which are stored within a data-store, can be retrieved in accordance with specific search parameters.

Trigger Events

The Audit Consumer sends a Retrieve ATNA Audit Events message when it needs to process or analyze CA:Aud audit records.

Message Semantics

The Retrieve ATNA Audit Event message shall be an HTTP GET request sent to the Audit Record Repository. This message is a FHIR search (see http://hl7.org/fhir/R4/search.html) on AuditEvent Resources (see http://hl7.org/fhir/R4/auditevent.html).

This “search” target is formatted as:

<scheme>://<authority>/<path>/AuditEvent?date=ge[start-time]&date=le[stop-time]&<query>

where:

- <scheme> shall be either http or https. The use of http or https is a policy decision, but https is usually appropriate due to confidentiality of CA:Aud audit record content.
- <authority> shall be represented as a host (either IP address or DNS name) followed optionally by a colon and port number.
- The Audit Record Repository may use <path> to segregate the HTTP search service for AuditEvent implementation from other REST-based services.
- At least one date search parameter is required. See section Date Search Parameters.
- “&” is a conditional parameter that shall be present if the <query> parameter is present.
- <query>, if present, represents a series of encoded name-value pairs representing filters for the search. See section Additional Search Parameters.

Date Search Parameters

The date parameter shall be used to specify an upper and/or lower bound for the search. At least one date parameter shall be present. Two date parameters are recommended in every search by the Audit Consumer and shall be supported by the Audit Record Repository in order to avoid overloading the Audit Consumer. These parameters allow the Audit Consumer to specify the time frame of creation of audit records of interest and enable the Audit Consumer to constrain the number of audit records returned. The values for the date search parameters shall be in RFC3339 format.

Note: RFC3339 format is the format mandated by Syslog for time stamps and is a sub-set of the XML date-time data format used by FHIR.

For example, to search AuditEvent Resources created during the whole day of January 5, 2013:

http://example.com/ARRservice/AuditEvent?date=ge2013-01-05&date=le2013-01-05

The Audit Record Repository shall apply matching criteria to AuditEvent Resources characterized by AuditEvent.recorded field valued within the time frame specified in the Request message.

The Audit Record Repository shall apply other date matching criteria following rules defined by FHIR specification (http://hl7.org/fhir/R4/search.html).

Additional Search Parameters

The search parameters in this section may be supported by the Audit Consumer and shall be supported by the Audit Record Repository. These parameters can be used by the Audit Consumer to refine search requests.

The Audit Consumer shall encode all search parameters per RFC3986 “percent” encoding rules. Although FHIR allows unconstrained use of AND OR operators to make queries of unlimited complexity, this transaction constrains the queries allowed:

- Multiple search parameters shall only be combined using AND “&” operator.
- The OR “,” operator shall be used only within a single search parameter that has multiple values.

Additional search parameters are listed below:

- address is a parameter of string type. This parameter specifies the identifier of the network access point (NetworkAccessPointID) of the user device that creates the audit record (this could be a device id, IP address, or some other identifier associated with a device).

The value of this parameter shall contain the substring to match.

For example:

http://example.com/ARRservice/AuditEvent?date=ge2013-01-01&date=le2013-01-02&address=192.168.0.1

The Audit Record Repository shall match this parameter with the AuditEvent.agent.network.address.

- agent.identifier is a parameter of token type. This parameter identifies the user that participated in the event that originates the audit record.

For example, to search AuditEvent Resources related to the user “admin”:

http://example.com/ARRservice/AuditEvent?date=ge2013-01-01&date=le2013-01-02&agent.identifier=admin

The Audit Record Repository shall match this parameter with the AuditEvent.agent.who.identifier field.

If a patient identifier it is used, the Audit Record Repository will return only the audit records where the patient is involved in the event as a user.

- patient.identifier is a parameter of token type. This parameter specifies the identifier of the patient involved in the event as a participant or as a user. The value of this parameter can contain the namespace URI (that represents the assigning authority for the identifier) and the identifier.

For example:

http://example.com/ARRservice/AuditEvent?date=ge2013-01-01&date=le2013-01-02&patient.identifier=urn:oid:1.2.3.4|5678

The Audit Record Repository shall match this parameter with the AuditEvent.agent.who.identifier and AuditEvent.entity.what.identifier where the reference resolve to a Patient.

- entity.identifier is a parameter of token type. This parameter specifies unique identifier for the object. The parameter value should be identified in accordance to the entity type.

For example:

- ?entity.identifier=urn:oid:1.2.3.4.5|123-203-FJ
- ?entity.identifier=|123-203-FJ.

The Audit Record Repository shall match this parameter with the AuditEvent.entity.what.identifier field that is of type identifier. If a patient identifier is used the Audit Record Repository will return only audit records where the patient is involved in the event as a participant.

- entity-type is a parameter of token type. This parameter specifies the type of the object (e.g., Person, System Object, etc.). The parameter value shall contain the namespace URI http://hl7.org/fhir/audit-entity-type or http://hl7.org/fhir/resource-types defined by FHIR and a coded value. See http://hl7.org/fhir/R4/valueset-audit-entitytype.html for codes that shall be used.

The Audit Record Repository shall match this parameter with the AuditEvent.entity.type field.

- entity-role is a parameter of token type. This parameter specifies the role played by the entity (e.g., Report, Location, Query, etc.). The parameter value shall contain the namespace URI http://hl7.org/fhir/object-role defined by FHIR and a coded value. See http://hl7.org/fhir/R4/object-role for codes that shall be used.

For example, to search all the audit records related to the document entity (Report=”3”) with the unique id 12345^1.2.3.4.5 a fully specified request would be: http://example.com/ARRservice/AuditEvent?date=ge2013-01-01&date=le2013-01-02&entity-role=http://hl7.org/fhir/object-role|3&entity-id=urn:oid:1.2.3.4.5|12345

The Audit Record Repository shall match this parameter with the AuditEvent.entity.role field.

- source.identifier is a parameter of token type. This parameter identifies the source of the audit event.

For example, to search AuditEvent Resources produced by the audit source application characterized by unique ID: 1234:

http://example.com/ARRservice/AuditEvent?date=ge2013-01-01&date=le2013-01-02&source=1234

The Audit Record Repository shall match this parameter with the AuditEvent.source.observer.identifier field.

- type is a parameter of token type. This parameter represents the identifier of the specific type of event audited. The parameter value shall contain the namespace URI http://dicom.nema.org/resources/ontology/DCM and a coded value. Codes available are defined by IHE (see Record Audit Event [ITI-20] section Send Audit Resource Request Message, Trigger Events.

For example, to search AuditEvent Resources related to PHI Export Events:

http://example.com/ARRservice/AuditEvent?date=ge2013-01-01&date=le2013-01-02&type=http://dicom.nema.org/resources/ontology/DCM|110106

The Audit Record Repository shall match this parameter with the AuditEvent.type field.

- subtype is parameter of token type. This parameter identifies the specific IHE transaction that originates the audit record. The parameter value can contain the namespace URI urn:ihe:event-type-code to search for audit messages triggered by IHE transactions with the defined audit message. Each IHE transaction that defines an CA:Aud messages, specifies a code identifying the transaction itself, and assigns this code to the EventTypeCode element within the [ITI-20] audit record.

For example, to search AuditEvent Resources related to Retrieve Document Set [ITI-43] transactions:

http://example.com/ARRservice/AuditEvent?date=ge2013-01-01&date=le2013-01-02&subtype=urn:ihe:event-type-code|ITI-43

The Audit Record Repository shall match this parameter with the AuditEvent.subtype field.

- outcome is a parameter of token type. This parameter represents whether the event succeeded or failed. The parameter value shall contain the namespace URI http://hl7.org/fhir/audit-event-outcome and a code taken from the related value set. See http://hl7.org/fhir/R4/valueset-audit-event-outcome.html.

To search AuditEvent Resources related to failed events:

http://example.com/ARRservice/AuditEvent?date=ge2013-01-01&date=le2013-01-02&outcome=http://hl7.org/fhir/audit-event-outcome|4,8,12

The Audit Record Repository shall match this parameter with the AuditEvent.outcome field.

The HL7® FHIR® standard provides additional search parameters. This transaction does not define specific behavior on those parameters (such as _sort, _include, etc.). See http://hl7.org/fhir/R4/search.html for details about available parameters.

Populating Expected Response Format

The HL7® FHIR® standard provides encodings for responses as either XML or JSON. The Audit Record Repository shall support both message encodings. The Audit Consumer shall support one and may optionally support both encodings. For Desired Response Encoding and format negotiation see ITI TF-2: Appendix Z.6.

Expected Actions

The Audit Record Repository (ARR) maintains a database of audit events. The Audit Record Repository retains data according to local policies, and some data may be deleted.

The Audit Record Repository shall return all the audit events stored in its database that match the query parameters, and which the requester is authorized to view (see CA:Aud Overview for further details).

When performing matching based on the search parameters, the Audit Record Repository shall:

- Select all audit records that have a time interval specified in the request URL.
- If search parameters other than those defined in section Additional Search Parameters (e.g., _sort, _include FHIR search result parameters) are specified in the request URL, then

o If the Audit Record Repository does not support the parameter, it shall be ignored;

o If the Audit Record Repository supports the parameter, the matching or other behavior shall comply with the matching rules for its datatype in FHIR.

The Audit Record Repository shall return matching resources using the Retrieve ATNA Audit Event Response Message. See section Retrieve ATNA Audit Event Response Message.

Retrieve ATNA Audit Event Response Message

The Audit Record Repository sends the Retrieve ATNA Audit Event Response message in response to a query from an Audit Consumer

Trigger Events

The Audit Record Repository creates this message when it receives and processes a Retrieve ATNA Audit Event message.

Message Semantics

When the Audit Record Repository successfully processes the search request, it shall return the matching AuditEvent Resources inside a FHIR Bundle Resource.

The “Content-Type” of the response will depend upon the response format negotiation described in ITI TF-2: Appendix Z.6.

If the date search parameter is missing (see section Date Search Parameters), the Audit Record Repository may return HTTP response code 400 - Bad Request.

If the specified search parameters do not result in any matching audit record, the Audit Record Repository shall return HTTP response of success 200, with an empty FHIR bundle.

If the requested data size is considered excessive by the Audit Record Repository, it may choose to return the results in a series of pages (see https://www.hl7.org/fhir/R4/http.html#paging).

Other HTTP response codes may be returned by the Audit Record Repository. See ITI TF-2: Appendix Z.7 Guidance on Access Denied Results.

The Audit Record Repository should complement the returned error code with a human readable description of the error condition.

Audit Record Repository may return HTTP redirect responses (responses with values of 301, 302, 303, or 307) in response to a request. Audit Consumers must follow redirects, but if a loop is detected, it may report an error.

FHIR Bundle of Audit Events Messages

When the search is successful, the body of the Response message shall contain a FHIR Bundle of AuditEvent Resources.

Example XML format:

```
<Bundle>
    <type>searchset</type>
    <total>3</total>
    <link>
        <relation value="self"/>
        <url value="http://example.com/ARRservice/AuditEvent?
date=&gt;2013-01-01&date=&lt;2013-01-02"/>
    </link>
    <entry>
        <fullUrl value="http://example.com/ARRservice/AuditEvent/23#"/>
        <resource> <AuditEvent> .....  </AuditEvent> </resource>
    </entry>
    <entry>
        <fullUrl value="http://example.com/ARRservice/AuditEvent/564#"/>
        <resource> <AuditEvent> ..... </AuditEvent> </resource>
    </entry>
    <entry>
        <fullUrl value="http://example.com/ARRservice/AuditEvent/3446#"/>
        <resource> <AuditEvent>  ..... </AuditEvent> </resource>
    </entry>
</Bundle>
```

Expected Actions

The Audit Consumer may further analyze the data received within the FHIR Bundle of AuditEvent Resources.

Security Considerations

See the general Security Considerations.

Audit Considerations

This transaction may involve the disclosure of sensitive information. Logging these retrieval transactions as a query event is appropriate.

However, CA:Aud does not require the Audit Record Repository to be able to send audit records using the Record Audit Event [ITI-20] transaction.

The following notation is used for optionality:

M	This field is mandatory.
U	The optionality of this field is unspecialized. The optionality of the underlying standard applies.
C	This field is mandatory if a specified condition is true.

The Audit Record Repository shall create and store locally an audit event as follows:

	Field Name	**Opt**	**Value Constraints**

Event AuditMessage/ EventIdentification	EventID	M	EV (110101, DCM, "Audit Log Used")
	EventActionCode	M	“R” (Read)
	EventDateTime	*U*	*not specialized*
	EventOutcomeIndicator	*U*	*not specialized*
	EventTypeCode	M	EV(“ITI-81”, “IHE Transactions”, “Retrieve ATNA AuditEvent”)
Source (Document Administrator) (1)			
Human Requestor (0..1)			
Destination (Document Registry) (1)			
Audit Source (Document Administrator) (1)			
AuditEvent Message (0..n)			

Where:

Source AuditMessage/ ActiveParticipant	*UserID*	*U*	*not specialized*
	AlternativeUserID	M	The process ID as used within the local operating system in the local system logs.
	UserName	*U*	*not specialized*
	UserIsRequestor	*U*	*not specialized*
	RoleIDCode	M	EV(110153, DCM, “Source”)
	NetworkAccessPointTypeCode	M	“1” for machine (DNS) name, “2” for IP address
	NetworkAccessPointID	M	The machine name or IP address.

Human Requestor (if known) AuditMessage/ ActiveParticipant	UserID	M	Identity of the human that initiated the transaction.
	AlternativeUserID	*U*	*not specialized*
	UserName	*U*	*not specialized*
	UserIsRequestor	*U*	*not specialized*
	RoleIDCode	M	Access Control role(s) the user holds that allows this transaction.
	NetworkAccessPointTypeCode	*U*	*not specialized*
	NetworkAccessPointID	*U*	*not specialized*
Destination AuditMessage/ ActiveParticipant	UserID	M	SOAP endpoint URI.
	AlternativeUserID	*U*	*not specialized*
	UserName	*U*	*not specialized*
	UserIsRequestor	*U*	*not specialized*
	RoleIDCode	M	EV(110152, DCM, “Destination”)
	NetworkAccessPointTypeCode	M	“1” for machine (DNS) name, “2” for IP address
	NetworkAccessPointID	M	The machine name or IP address.
Audit Source AuditMessage/ AuditSourceIdentification	*AuditSourceID*	*U*	*not specialized*
	AuditEnterpriseSiteID	*U*	*not specialized*
	AuditSourceTypeCode	*U*	*not specialized*

AuditEvent Message AuditMessage/ ParticipantObjectIdentification	ParticipantObjectTypeCode	M	“2” (System object)
	ParticipantObjectTypeCode Role	M	“13” (Security Resource)
	ParticipantObjectDataLifeCycle	*U*	*not specialized*
	ParticipantObjectIDTypeCode	M	EV(“12”, “RFC-3881”, "URI")
	ParticipantObjectSensitivity	*U*	*not specialized*
	ParticipantObjectID	M	The URI of the Audit log
	ParticipantObjectName	M	“Security Audit Log”
	ParticipantObjectQuery	*U*	*not specialized*
	ParticipantObjectDetail	*U*	*not specialized*

1.3.2 Internet User Authorization (IUA)

Overview

The Internet User Authorization (IUA) is an interoperability profile that provides an authorization profile for the HTTP RESTful transactions. Being authorized means that the user, patient, or provider has legitimate access to this HTTP RESTful service. The authorization includes identifying the user and the application that is making the request to the HTTP RESTful server, so that server can make further access control decisions.

IUA conveys User Identity, Attributes, and Authorizations to a RESTful service to enable security and confidentiality policy enforcement. The primary use cases are for obtaining authorization for access to a resource using HTTP RESTful HTTP transactions. There are other use cases for delegation, provisioning, etc. which are out of scope for this profile.

The authorization service is separated from the HTTP RESTful access so that it can be provided by a different organization or part of the organization than the resource service. This is driven by the requirements of patients, providers, and other users to simplify and maintain autonomy and control over authorization services. A user may interact with dozens of providers. It is difficult for the user to coordinate different authorization mechanisms with each of these dozens of providers.

This pattern is a common Internet usage and there are already vendors of authorization services that are being used to solve this problem. These include Facebook, Google, and a variety of other service providers in different commercial and governmental sectors. Some countries may use their citizen identity card to access their governmental services. These overlap with providers of authentication services. These services allow a patient to establish an authentication and authorization relationship with minimal provisioning by the healthcare provider. The user can specify “use vendor X” to their healthcare provider.

Actors and Transactions

The following diagram provides an overview of the IUA profile Actors, Transactions and their interactions.

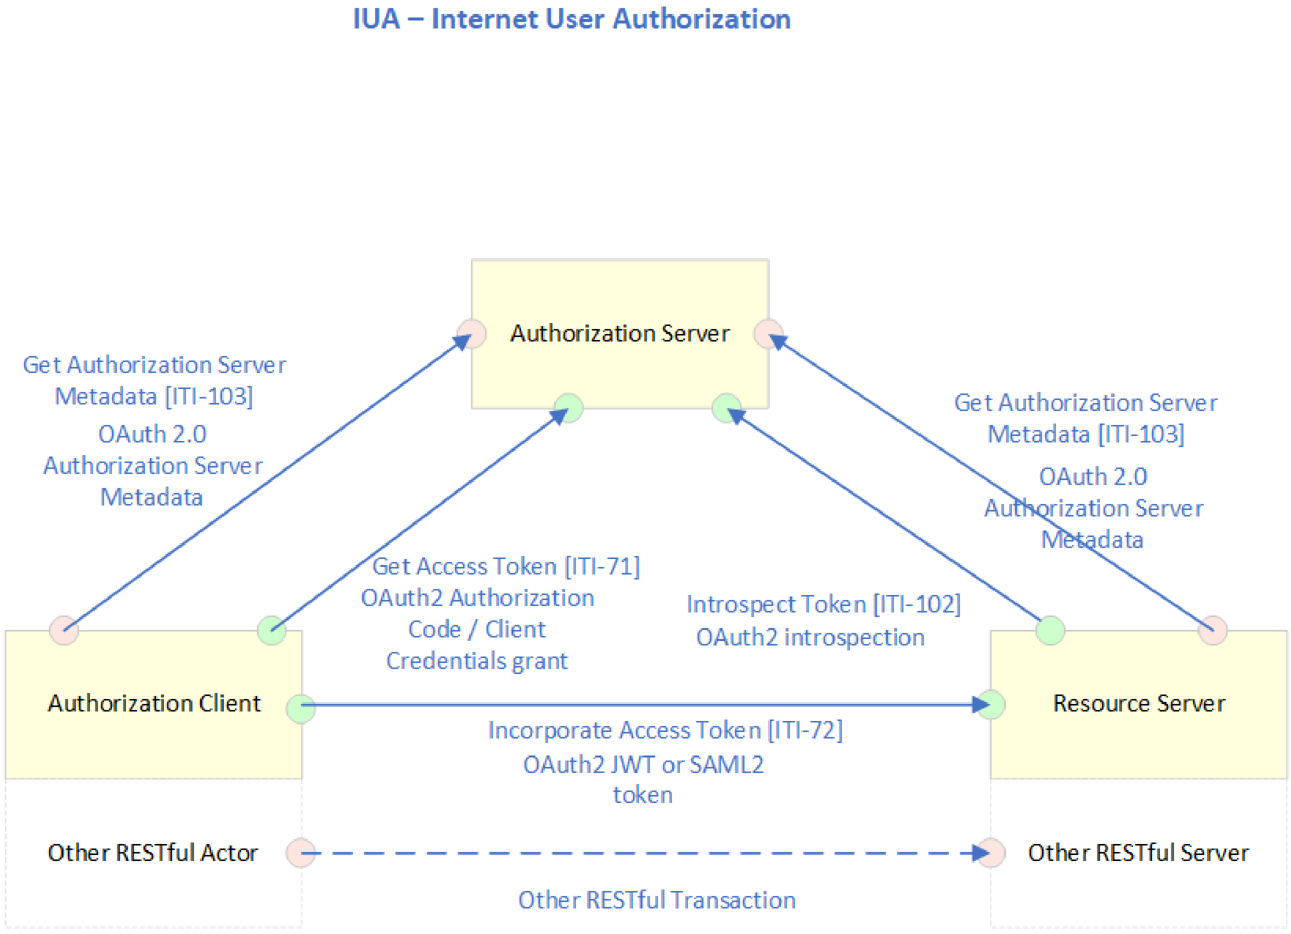
IUA – Internet User Authorization
Authorization Server
Get Authorization Server Metadata [ITI-103]
OAuth 2.0 Authorization Server Metadata
Get Authorization Server Metadata [ITI-103]
OAuth 2.0 Authorization Server Metadata
Get Access Token [ITI-71]
OAuth2 Authorization Code / Client Credentials grant
Introspect Token [ITI-102]
OAuth2 introspection
Authorization Client
Resource Server
Incorporate Access Token [ITI-72]
OAuth2 JWT or SAML2 token
Other RESTful Actor
Other RESTful Server
Other RESTful Transaction

The table below lists the transactions for each actor directly involved in the IUA profile. To claim compliance with IUA, an actor shall support all required transactions (labeled “R”) and may support the optional transactions (labeled “O”).

Actor	Transaction	Optionality
Authorization Client	Get Access Token [ITI-71] Incorporate Access Token [ITI-72] Get Authorization Server Metadata [ITI-103]	R R O
Authorization Server	Get Access Token [ITI-71] Get Authorization Server Metadata [ITI-103] Introspect Token [ITI-102]	R O O
Resource Server	Incorporate Access Token [ITI-72] Get Authorization Server Metadata [ITI-103] Introspect Token [ITI-102]	R O O

Transactions

- Get Access Token [ITI-71] - This transaction is used by an Authorization Client to retrieve an OAuth 2.1-compliant access token. Uses RFC6749 - OAuth2.1 protocol and RFC7519 – JSON Web Token and JSON Web Token (JWT) Profile for OAuth 2.0 Access Tokens.
- Incorporate Access Token [ITI-72] - This transaction is used to incorporate authorization information into HTTP RESTful transactions. Uses RFC6749 - OAuth2.1 protocol.
- Introspect Token [ITI-102] - Token introspection defines a protocol that allows Resource Servers to query the Authorization Server to determine the set of claims for a given token that was presented to them by an Authorization Client. These claims include whether the token is currently active and the authorization context in which the token was granted.

Sequence Diagram

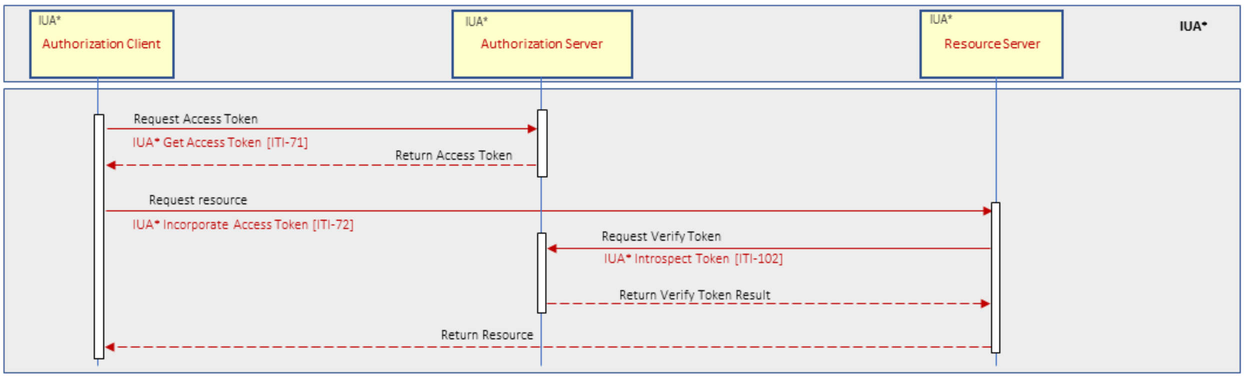

IUA - Canadian Implementation Guidance

The IUA - Canadian Implementation Guidance describes the recommendations for the IHE IUA (Internet User Authorization) profile.

IUA Actor Options

Actor	**Options**	**Notes**
Authorization Server	JWT Token	
	Token Introspection	(Note 1)
	Signature Verification	(Note 1)
	Authorization Server Metadata	
Resource Server	JWT Token	
	Token Introspection	(Note 1)
	Signature Verification	(Note 1)
	Authorization Server Metadata	
Authorization Client	Authorization Server Metadata	

Note 1: The actor shall support one of the two options: Token Introspection or Signature Verification (using Authorization Server's Public Key)

JWT Token Option

This option uses JSON Web Token encoding for access tokens issued by the Authorization Server.

JWT IUA Extension

The JWT IUA extension, described in the JWT IUA extension section of the IUA specification, are currently considered out of scope.

Token Introspection Option

The Token Introspection Option allows Resource Servers to check the origin and validity of the access token.

The Resource Server delegates the token verification to the Authorization Server by accessing the Token Introspection endpoint. The Authorization Server will perform all applicable checks against a token's state, such as checking whether the token has expired, verifying signatures, etc.

Also, generally the token introspection allows to find out additional information such as the user and scopes that are associated with the token, however in case of JWT tokens this is not necessary, as the resource server is able to parse the token and perform these additional verifications.

Signature Verification Option

This option allows Resource Servers to check the origin and validity of the JWT access token without a network call to the Authorization Server. The Resource Server is using the Authorization Server's public key to validate the signature of the token, and then parse the claims within the structured token itself.

The Authorization Server's public key endpoint can be discovered using the well-known URI (see Authorization Server Metadata Option). If the Authorization Server metadata option is not supported by the Authorization Server, the public key can be obtained by other means.

Authorization Server Metadata Option

Server endpoints are available to be discovered via the well-known URI discovery mechanism. The client issues a HTTP GET request to the well-known metadata endpoint associated with Authorization Server.

An example of such a request using [RFC8414] is:

BASE_URL/.well-known/openid-configuration

where BASE_URL includes realm (tenant) information.

Keycloak authorization server example:

- BASE_URL

  https://<keycloakserver>/realms/<realm>
- Authorization server metadata discovery endpoint

  BASE_URL/.well-known/openid-configuration
- Authorization server certificate endpoint (discoverable via well-known URI)

  BASE_URL/protocol/openid-connect/certs

Amazon Cognito authorization server example:

- BASE_URL

  https://cognito-idp.<region>.amazonaws.com/<user_pool_id>
- Authorization server metadata discovery endpoint

  BASE_URL/.well-known/openid-configuration
- Authorization server certificate endpoint (discoverable via well-known URI)

  BASE_URL/.well-known/jwks.json

Authorization Grant Types

Actor	**Options**	**Notes**
Authorization Server	Authorization Code	Support both
	Client Credentials	
Authorization Client	Authorization Code	Support one
	Client Credentials	

Security Considerations

Recommendations for Authorization Server and Authorization Client that are using the Authorization Code flow:

- PKCE (Proof Key for Code Exchange) – Mechanism to protect against CSRF (Cross-site request forgery) and Code Injection
- state – Request parameter that enables additional protection against CSRF
- nonce – Request parameter that enables enforce one-time code usage and protection against replay attacks

1.3.3 Consistent Time (CT)

Overview

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, which is sufficient for most purposes.

Various infrastructure, security, and acquisition profiles require use of a consistent time base on multiple computers, to synchronize logs, authenticate users, digitally sign documents, etc.

Actors and Transactions

The following diagram provides an overview of the CT profile Actors, Transactions and their interactions.

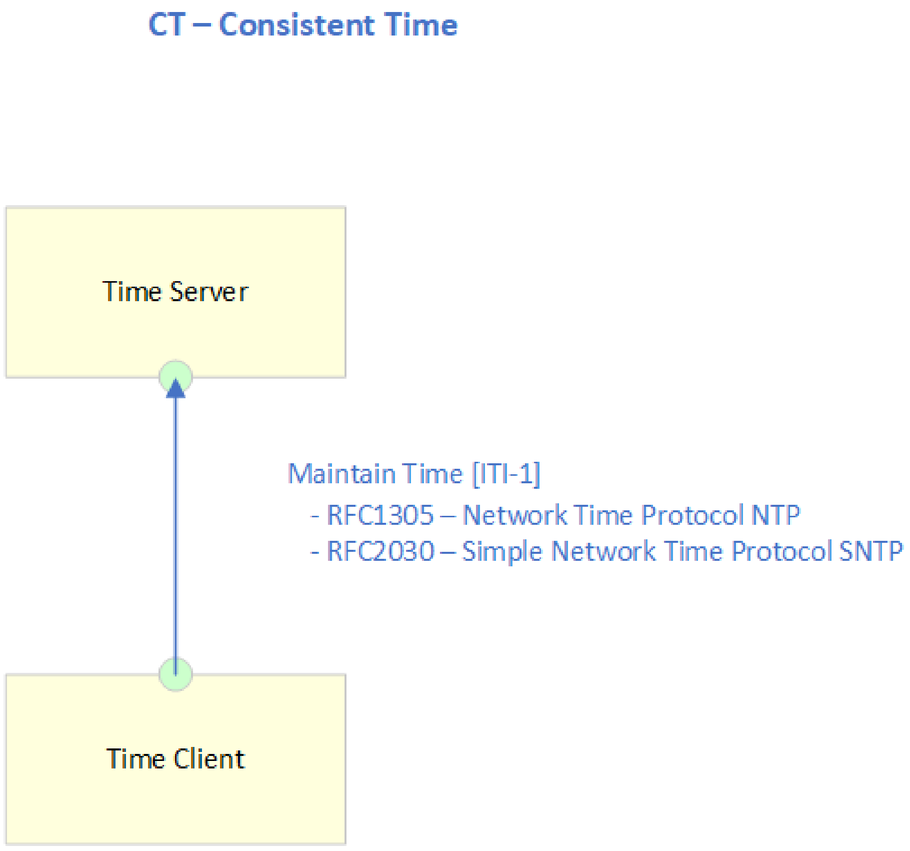
CT – Consistent Time
Time Server
Maintain Time [ITI-1]
- RFC1305 – Network Time Protocol NTP
- RFC2030 – Simple Network Time Protocol SNTP
Time Client

The table below lists the transactions for each actor directly involved in the CT profile. To claim compliance with CT, an actor shall support all required transactions (labeled “R”).

Actor	Transaction	Optionality
Time Client	Maintain Time [ITI-1]	R
Time Server	Maintain Time [ITI-1]	R

Transactions

- Maintain Time [ITI-1] - Uses NTP (RFC1305) or SNTP (RFC4330) time service responses to maintain synchronization with Time Servers and maintain the local system clock.

Sequence Diagram

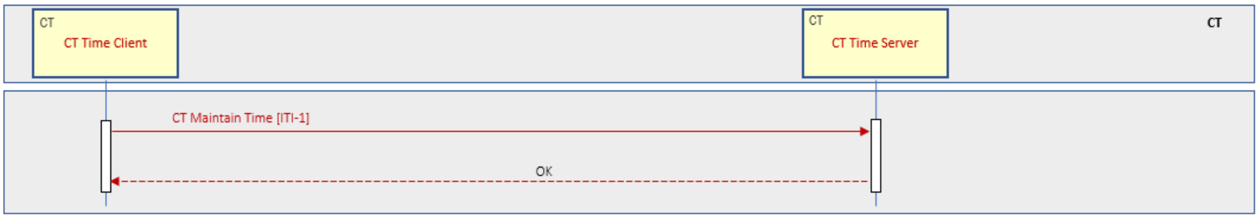

CT - Canadian Implementation Guidance

The CT - Canadian Implementation Guidance describes the recommendations for the IHE CT (Canadian Consistent Time) profile.

Time Server

The Time Server shall synchronize to stratum-2 time servers or better.

In Canada, the National Research Council provides NTP services, including Secure NTP, on the NRC stratum-2 servers free of charge. See NRC for more information.

Examples of NRC NTP time servers:

- nrc.ca
- chu.nrc.ca

1.3.4 Sharing Valuesets, Codes and Maps (SVCM)

Overview

The Sharing ValueSets, Codes and Maps (SVCM) Profile defines a lightweight interface through which healthcare systems may retrieve centrally managed uniform nomenclature and mappings between code systems based on the HL7 Fast Healthcare Interoperability Resources (FHIR) specification.

SVCM supports querying for value sets and code systems using the standard HL7 FHIR resources. It also supports looking up and validating codes as well as expanding a value set to list all the available codes.

Optionally concept maps can also be included to translate from one code system or value set to another (e.g. SNOMED CT to LOINC).

Terminologies managed in value sets are most useful when they are widely shared and standardized across geography and disciplines to add clarity and specificity.

Actors and Transactions

The following diagram provides an overview of the SVCM profile Actors, Transactions and their interactions.

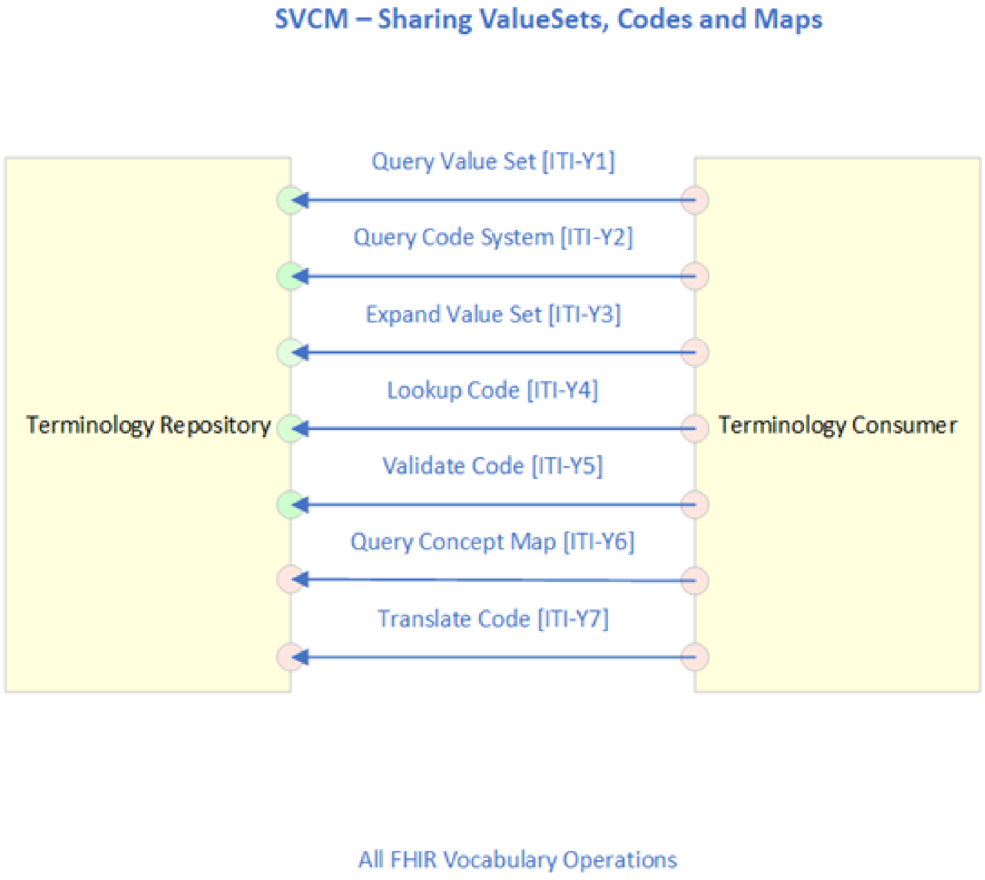
SVCM – Sharing ValueSets, Codes and Maps
Terminology Repository
Terminology Consumer
Query Value Set [ITI-Y1]
Query Code System [ITI-Y2]
Expand Value Set [ITI-Y3]
Lookup Code [ITI-Y4]
Validate Code [ITI-Y5]
Query Concept Map [ITI-Y6]
Translate Code [ITI-Y7]
All FHIR Vocabulary Operations

The table below lists the transactions for each actor directly involved in the SVCM profile. To claim compliance with SVCM, an actor shall support all required transactions (labeled "R") and may support the optional transactions (labeled "O").

Actor	Transaction	Optionality
Terminology Consumer	Query Value Set [ITI-95]	O (*)
	Query Code System [ITI-96]	O (*)
	Expand Value Set [ITI-97]	O (*)
	Lookup Code [ITI-98]	O (*)
	Validate Code [ITI-99]	O (*)
	Query Concept Map [ITI-100]	O
	Translate Code [ITI-101]	O

Actor	Transaction	Optionality
Terminology Repository	Query Value Set [ITI-95] Query Code System [ITI-96] Expand Value Set [ITI-97] Lookup Code [ITI-98] Validate Code [ITI-99] Query Concept Map [ITI-100] Translate Code [ITI-101]	R R R R R O O

(*) A Terminology Consumer shall support at least one of these transactions.

Transactions

- Query Value Set [ITI-95] - used by the Terminology Consumer to find value sets based on criteria it provides in the query parameters of the request message, or to retrieve a specific value set. Uses a FHIR ValueSet query.
- Query Code System [ITI-96] - used by the Terminology Consumer to solicit information about code systems whose data match data provided in the query parameters on the request message. Uses a FHIR CodeSystem query.
- Expand Value Set [ITI-97] - used by the Terminology Consumer to expand a given ValueSet to return the full list of concepts available in that ValueSet. Uses a FHIR $expand ValueSet operation.
- Lookup Code [ITI-98] - used by the Terminology Consumer to lookup a given code to return the full details. Uses a FHIR $lookup ValueSet or CodeSystem operation.
- Validate Code [ITI-99] - used by the Terminology Consumer to validate the existence of a given code in a value set or code system. Uses a FHIR $validate-code ValueSet or CodeSystem operation.
- Query Concept Map [ITI-100] - used by the Terminology Consumer that supports the Translate Option to solicit information about concept maps whose data match data provided in the query parameters on the request message. Uses a FHIR ConceptMap query.
- Translate Code [ITI-101] - used by the Terminology Consumer that supports the Translate Option to translate a given code from a ValueSet to a code from another ValueSet based on a ConceptMap Resource. Uses a FHIR $translate ConceptMap operation.

Sequence Diagram

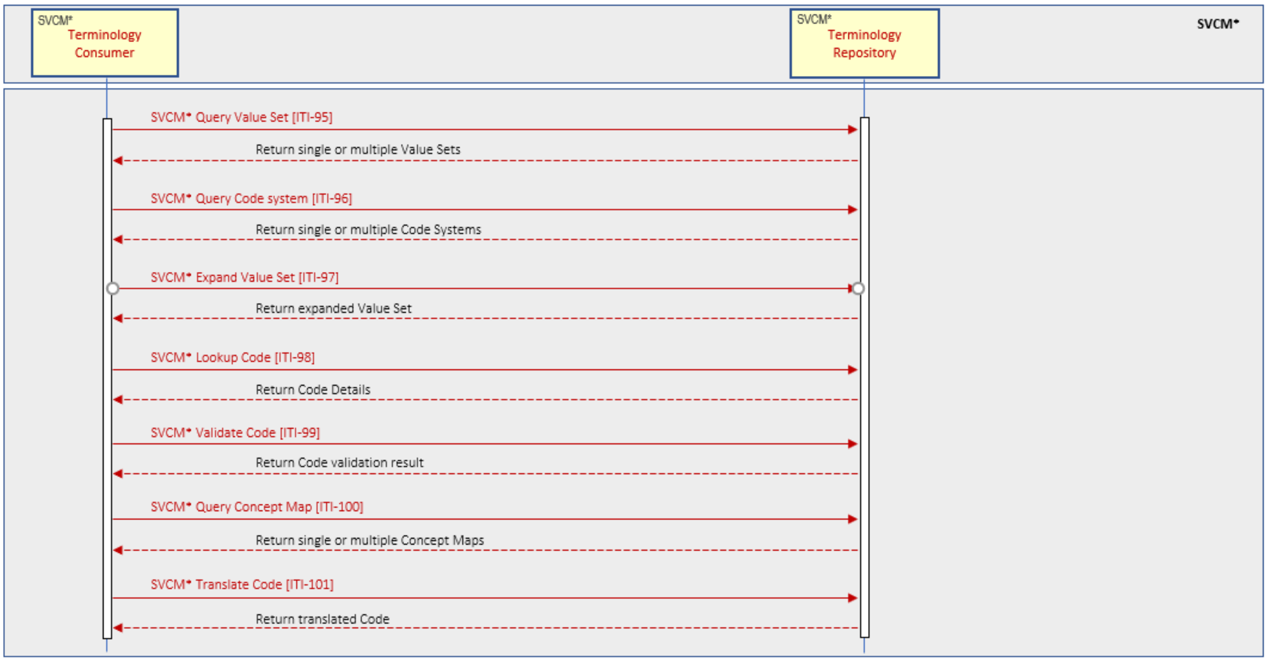
SVCM*
Terminology Consumer
SVCM*
Terminology Repository
SVCM*
SVCM* Query Value Set [ITI-95]
Return single or multiple Value Sets
SVCM* Query Code system [ITI-96]
Return single or multiple Code Systems
SVCM* Expand Value Set [ITI-97]
Return expanded Value Set
SVCM* Lookup Code [ITI-98]
Return Code Details
SVCM* Validate Code [ITI-99]
Return Code validation result
SVCM* Query Concept Map [ITI-100]
Return single or multiple Concept Maps
SVCM* Translate Code [ITI-101]
Return translated Code

Terminology Gateway

Value Sets that are part of pan-Canadian specifications are published in the Terminology Gateway. The Terminology Gateway provides a FHIR API interface compatible with the Terminology Repository SVCM actor, and implements the following transactions:

- Query Value Set [ITI-95]
  - e.g. GET https://fhir.infoway-inforoute.ca/ValueSet?name=*route*
- Expand Value Set [ITI-97]
  - e.g. GET https://fhir.infoway-inforoute.ca/ValueSet/routeofadministration/$expand
- Validate Code [ITI-99]
  - e.g. GET https://fhir.infoway-inforoute.ca/ValueSet/routeofadministration/$validate-code?code=697971008&system=http%3A%2F%2Fsnomed.info%2Fsct
- Query Concept Map [ITI-100]
  - e.g. GET https://fhir.infoway-inforoute.ca/ConceptMap?name=*map*
- Translate Code [ITI-101]
  - e.g. GET https://fhir.infoway-inforoute.ca/ConceptMap/MP-NTP-Mapping/$translate?code=00000817

For performance reasons, the Terminology Gateway isn't supposed to be queried at runtime. Instead, it is recommended to rely on its notification services to update a local database to be used for validation.

1.3.5 Canadian Formatting Service (CA:FMT)

Introduction

CA:FMT is a Canadian Integration Specification that provides formatting support service. It provides support for transformation of documents between different formats (e.g. from FHIR to PDF, CDA, etc.).

Actors and Transactions

The following diagram provides an overview of the CA:FMT Actors, Transactions and their interactions.

CA:FMT – Canadian Formatting Service

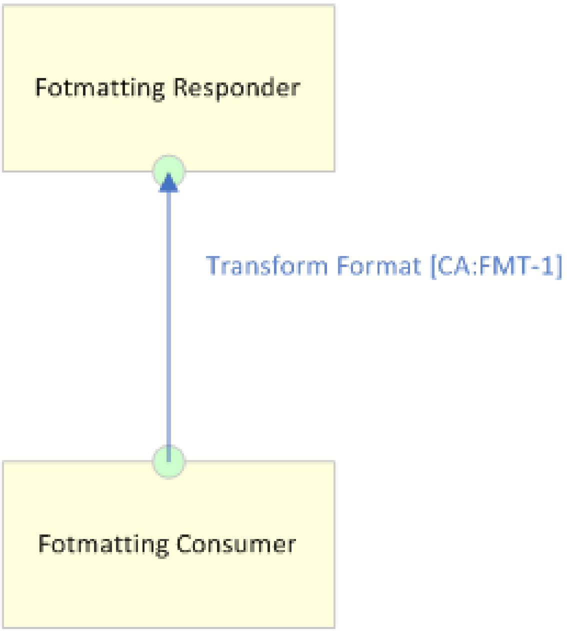

The table below lists the transactions for each actor directly involved in CA:FMT. To claim compliance with CA:FMT, an actor shall support all required transactions (labeled “R”).

Actors	Transactions	Optionality
Formatting Consumer	Transform Format [CA:FMT-1]	R
Formatting Responder	Transform Format [CA:FMT-1]	R

Transactions

- Transform Format [CA:FMT-1]: This transaction is used to transform a document from one format to another (e.g. FHIR to PFD, CDA, etc.)

Sequence Diagram

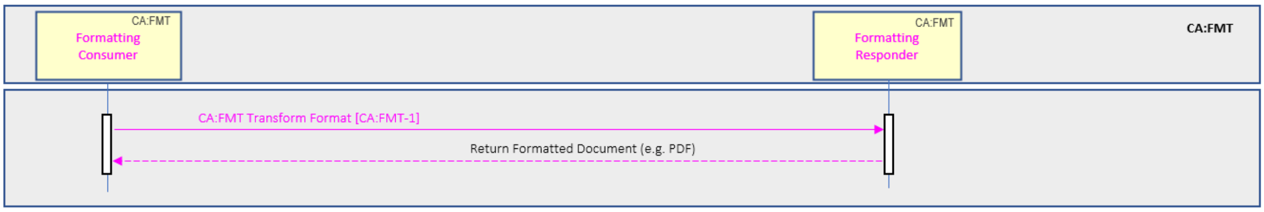
CA:FMT
Formatting Consumer
CA:FMT
Formatting Responder
CA:FMT
CA:FMT Transform Format [CA:FMT-1]
Return Formatted Document (e.g. PDF)

1.4 Document Exchange Profiles

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

Overview

The pan-Canadian FHIR Exchange (CA:FeX) is an implementable, testable interoperability specification based on HL7 FHIR Implementation Guides, that defines building blocks to enable creating, consuming and sharing clinical data via FHIR RESTful exchange patterns.

For details, refer to the pan-Canadian FHIR Exchange Interoperability Specifications (CA:FeX v2.1.0 DRAFT-preBallot).

The following are some examples of the benefits of CA:FeX:

- Supports scenarios of health care provider creating, viewing and updating documents using standardized HL7 FHIR operations such as submission, search and retrieval
- Supports safe provision of care in a scheduled or unscheduled medical situation
- Supports transitions of care or transfers of patients across the continuum of care
- Supports coordination and collaboration of a patient’s care

Actors and Actor Options

CA:FeX actor options are organized into four options providing implementers with choices that most suit their context and needs, while minimizing implementation details not relevant to them. Implementers may choose to opt into one or more actor option. The table below summarizes the CA:FeX actor options.

Actor Option	**Description**
A. Bundle Option	Exchange of data using /Bundle endpoint
B. Metadata Option	Exchange of data through metadata resources

C. Single Resource Option	Exchange of data using FHIR single resource endpoints
D. Summary Option	Exchange of document using $summary operation

Transaction Naming Convention

CA:FeX transactions aim to follow a consistent numbering convention for specific types of data exchange. The table below summarizes the numbering conventions commonly used for different types of transactions (see details in the next section "Actors and Transactions")

Transaction Number	**Description**	**Example**
CA:FeX-1	Submit Data	Submit FHIR Document (CA:FeX-1A) Generate Metadata (CA:FeX-1B)
CA:FeX-2	Search Data	Search FHIR Document (CA:FeX-2A) Fetch DocumentReference (CA:FeX-2B)
CA:FeX-3	Retrieve Data	Retrieve FHIR Document (CA:FeX-3A) Retrieve Document (CA:FeX-3B)

For example, CA:FeX-1A likely involves a Submit Data action and CA:FeX-2B involves a Search Data action.

Actors and Transactions

The following diagram provides an overview of the CA:FeX Actors, Transactions and their interactions.

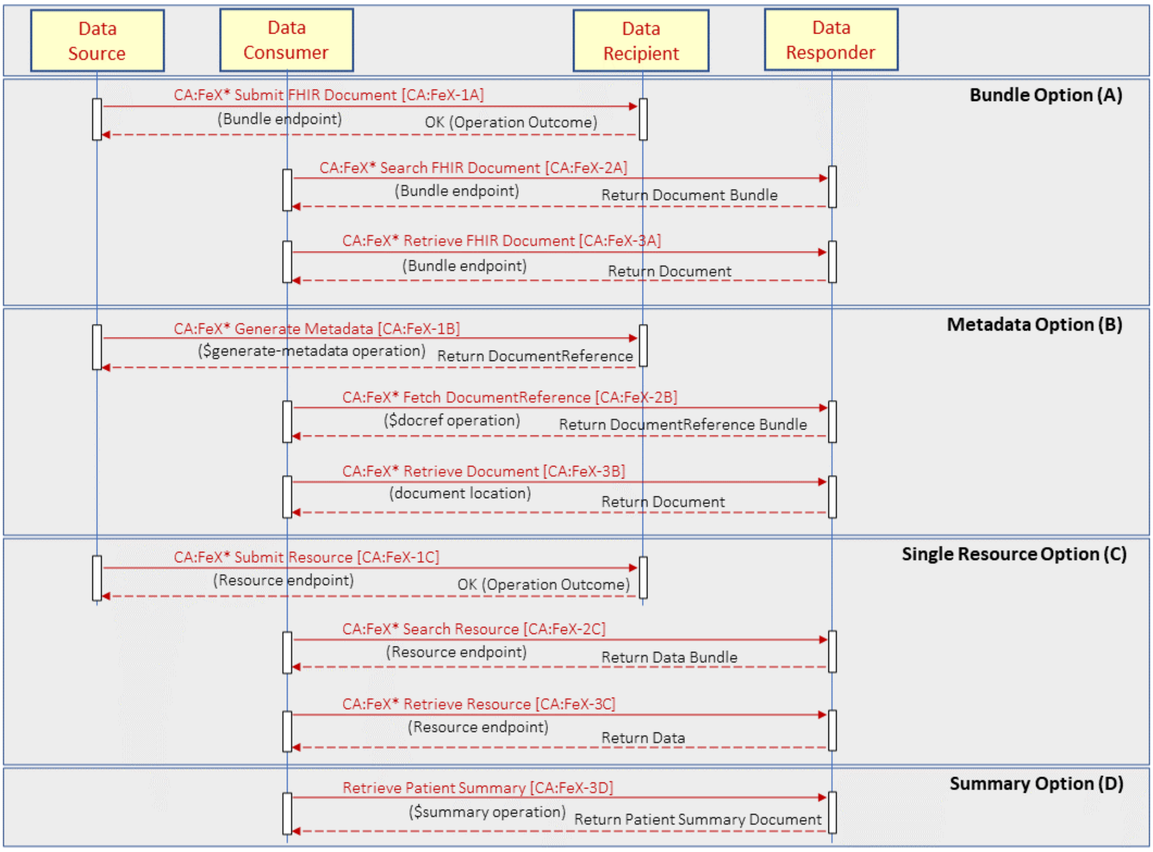
Data Source
Data Consumer
Data Recipient
Data Responder
Bundle Option (A)
CA:FeX* Submit FHIR Document [CA:FeX-1A]
(Bundle endpoint)
OK (Operation Outcome)
CA:FeX* Search FHIR Document [CA:FeX-2A]
(Bundle endpoint)
Return Document Bundle
CA:FeX* Retrieve FHIR Document [CA:FeX-3A]
(Bundle endpoint)
Return Document
Metadata Option (B)
CA:FeX* Generate Metadata [CA:FeX-1B]
($generate-metadata operation)
Return DocumentReference
CA:FeX* Fetch DocumentReference [CA:FeX-2B]
($docref operation)
Return DocumentReference Bundle
CA:FeX* Retrieve Document [CA:FeX-3B]
(document location)
Return Document
Single Resource Option (C)
CA:FeX* Submit Resource [CA:FeX-1C]
(Resource endpoint)
OK (Operation Outcome)
CA:FeX* Search Resource [CA:FeX-2C]
(Resource endpoint)
Return Data Bundle
CA:FeX* Retrieve Resource [CA:FeX-3C]
(Resource endpoint)
Return Data
Summary Option (D)
Retrieve Patient Summary [CA:FeX-3D]
($summary operation)
Return Patient Summary Document

*FHIR Enabled

The tables below list the transactions for each actor directly involved in CA:FeX. To claim compliance with CA:FeX, an actor shall support all required transactions (labeled “R”).

Bundle Option (A)

Actor	Transaction	Endpoint Info (Summary)	Optionality
Data Source	Submit FHIR Document [CA:FeX-1A]	/Bundle	R
Data Consumer	Search FHIR Document [CA:FeX-2A]	/Bundle	R
	Retrieve FHIR Document [CA:FeX-3A]	/Bundle	R
Data Recipient	Submit FHIR Document [CA:FeX-1A]	/Bundle	R

Data Responder	Search FHIR Document [CA:FeX-2A]	/Bundle	R
	Retrieve FHIR Document [CA:FeX-3A]	/Bundle	R

- Submit FHIR Document [CA:FeX-1A] – This transaction performs a FHIR Create interaction to transfer a FHIR Document, typically Bundle of type Composition.
- Search FHIR Document [CA:FeX-2A] – This transaction performs a FHIR Search interaction to obtain FHIR Documents that satisfies a set of parameters.
- Retrieve FHIR Document [CA:FeX-3A] – This transaction performs a FHIR Read interaction to retrieve a FHIR Document.

Metadata Option (B)

Actor	**Transaction**	**Endpoint Info (Summary)**	**Optionality**	**Notes**
Data Source	Generate Metadata [CA:FeX-1B]	$generate-metadata	R	Based on MHD
Data Consumer	Fetch DocumentReference [CA:FeX-2B]	$docref	R	Based on IPA
	Retrieve Document [CA:FeX-3B]	/Bundle, /Binary	R	
Data Recipient	Generate Metadata [CA:FeX-1B]	$generate-metadata	R	Based on MHD
Data Responder	Fetch DocumentReference [CA:FeX-2B]	$docref	R	Based on IPA
	Retrieve Document [CA:FeX-3B]	/Bundle, /Binary	R	

- Generate Metadata [CA:FeX-1B] – This transaction performs a $generate-metadata operation to submit a FHIR Document and provides a DocumentReference associated to the submitted document.
- Fetch DocumentReference [CA:FeX-2B] – This transaction performs $docRef operation to obtain DocumentReferences that contain links to FHIR Document or FHIR Binary resources from the Data Responder. Note that other options that could be possible are not recommended in CA:FeX (e.g. DocumentReference embedded documents).
- Retrieve Document [CA:FeX-3C] – This transaction uses /Bundle or /Binary (recommended) endpoints to retrieve data.

Single Resource Option (C)

Actor	Transaction	Endpoint Info (Summary)	Optionality
Data Source	Submit Resource [CA:FeX-1C]	/<type>	R
Data Consumer	Search Resource [CA:FeX-2C]	/<type>	R
	Retrieve Resource [CA:FeX-3C]	/<type>	R
Data Recipient	Submit Resource [CA:FeX-1C]	/<type>	R
Data Responder	Search Resource [CA:FeX-2C]	/<type>	R
	Retrieve Resource [CA:FeX-3C]	/<type>	R

- Submit Resource [CA:FeX-1C] – This transaction performs a FHIR Create interaction to submit a single FHIR resource.
- Search Resource [CA:FeX-2C] – This transaction performs a FHIR Search interaction to find FHIR resources matching a set of parameters.
- Retrieve Resource [CA:FeX-3C] – This interaction performs a FHIR Read to retrieve a FHIR resource using a known resource ID.

Summary Option (D)

Actor	Transaction	Endpoint Info (Summary)	Optionality	Notes
Data Consumer	Retrieve Patient Summary [CA:FeX-3D]	/Patient/$summary /Patient/[id]/$summary	R	Based on IPS
Data Responder	Retrieve Patient Summary [CA:FeX-3D]	/Patient/$summary /Patient/[id]/$summary	R	Based on IPS

- Retrieve Patient Summary [CA:FeX-3D] – This transaction uses the $summary operation on the Patient FHIR endpoint to retrieve a Patient Summary document.

1.4.2 Mobile Access to Health Documents (MHD)

Overview

The Mobile Access to Health Documents (MHD) Profile defines one standardized interface to health document sharing. This profile is applicable to systems where needs are simple, such as pulling the latest summary for display.

The following are examples of environments which may choose the MHD Profile:

- Medical devices such as those targeted by the Patient Care Devices (PCD) domain or Continua organization, submitting data in the form of documents.
- Kiosks used by patients in hospital registration departments, where it is anticipated that a hospital staff member will review, edit, and approve the document before it is allowed into the hospital system.
- PHR publishing into a staging area for subsequent import into an EHR or HIE.
- Patient or provider application that is configured to securely connect to a PHR in order to submit a medical history document. (For example BlueButton+)
- Electronic measurement device participating in an XDW workflow and pulling medical history documents from an HIE.
- A General Practitioner physician’s office with minimal IT capabilities using a mobile application to connect to an HIE or EHR

Actors and Transactions

The following diagram provides an overview of the MHD profile Actors, Transactions and their interactions.

MHD – Mobile Access to Health Documents

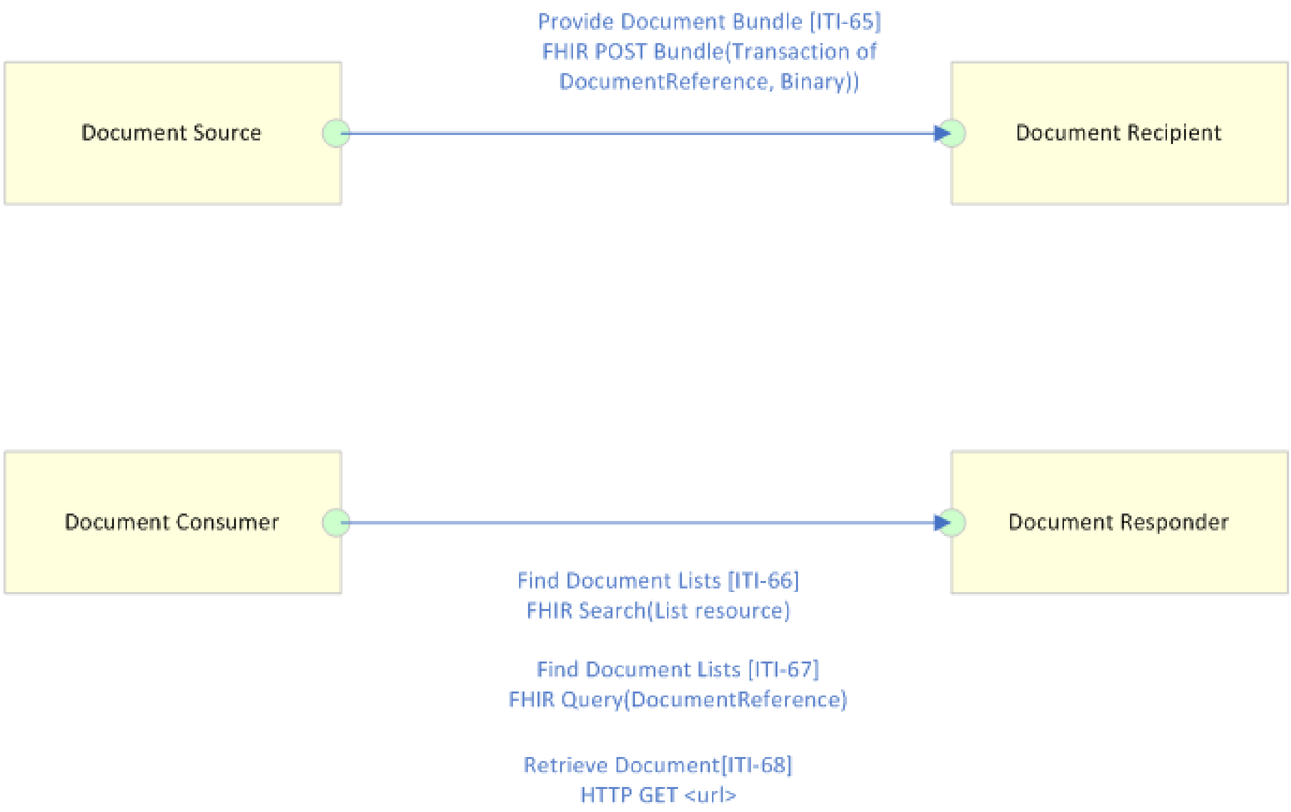
Provide Document Bundle [ITI-65]
FHIR POST Bundle(Transaction of
DocumentReference, Binary))
Document Source
Document Recipient
Document Consumer
Document Responder
Find Document Lists [ITI-66]
FHIR Search(List resource)
Find Document Lists [ITI-67]
FHIR Query(DocumentReference)
Retrieve Document[ITI-68]
HTTP GET <url>

The table below lists the transactions for each actor directly involved in the MHD profile. To claim compliance with MHD, an actor shall support all required transactions (labeled "R").

Actor	Transaction	Optionality
Document Source	Provide Document Bundle [ITI-65]	R
Document Recipient	Provide Document Bundle [ITI-65]	R
Document Consumer	Find Document Lists [ITI-66] Find Document References [ITI-67] Retrieve Document [ITI-68]	R R R

Actor	Transaction	Optionality
Document Responder	Find Document Lists [ITI-66] Find Document References [ITI-67] Retrieve Document [ITI-68]	R R R

Transactions

- Provide Document Bundle [ITI-65] – This transaction passes a Provide Document Bundle Request from a Document Source to a Document Recipient. Uses a FHIR POST for a *Bundle* of *DocumentReference* and/or *Binary* resources.
- Find Document Lists [ITI-66] - Used to locate and return metadata for previously stored document submission sets or folders. Uses a FHIR search returning a *List* of resources.
- Find Document References [ITI-67] - Used to find DocumentReference Resources that satisfy a set of query parameters. Uses a FHIR search that returns a Bundle of *DocumentReference* resources.
- Retrieve Document [ITI-68] - Used by the Document Consumer to retrieve a document from the Document Responder. Uses a FHIR GET operation returning a single *Document*.

Sequence Diagram

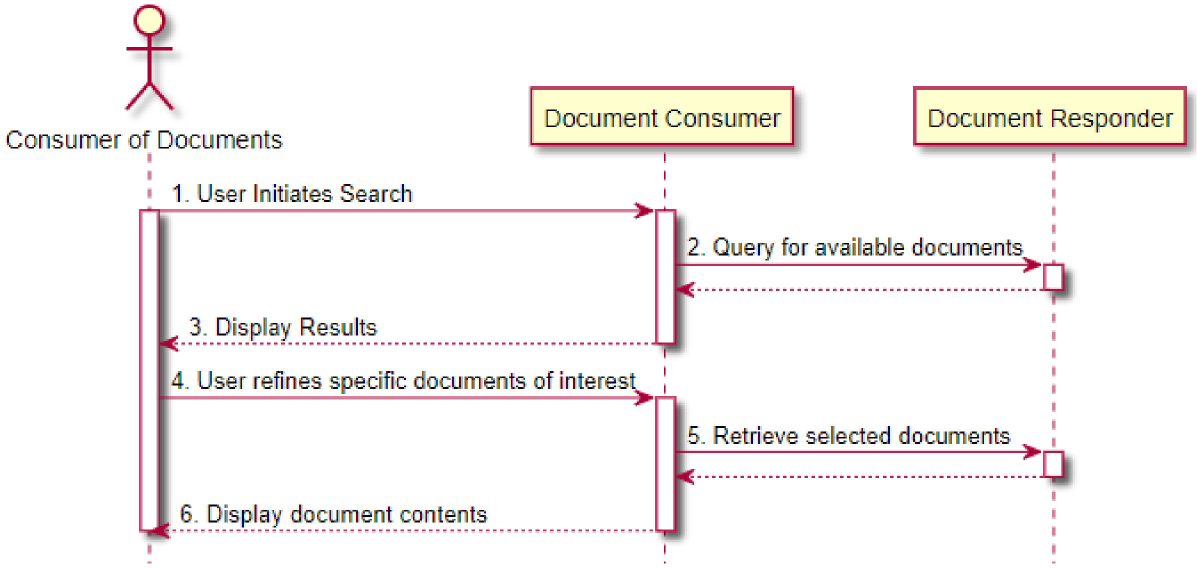
Consumer of Documents
Document Consumer
Document Responder
1. User Initiates Search
2. Query for available documents
3. Display Results
4. User refines specific documents of interest
5. Retrieve selected documents
6. Display document contents

1.4.3 Cross-Enterprise Document Media Interchange (XDM)

Overview

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.

XDM Facilitates person-to-person exchange of the healthcare information by supporting transport via physical media - USB and CD-R and supporting transport as an attachment to an email

Actors and Transactions

The following diagram provides an overview of the XDM profile Actors, Transactions and their interactions.

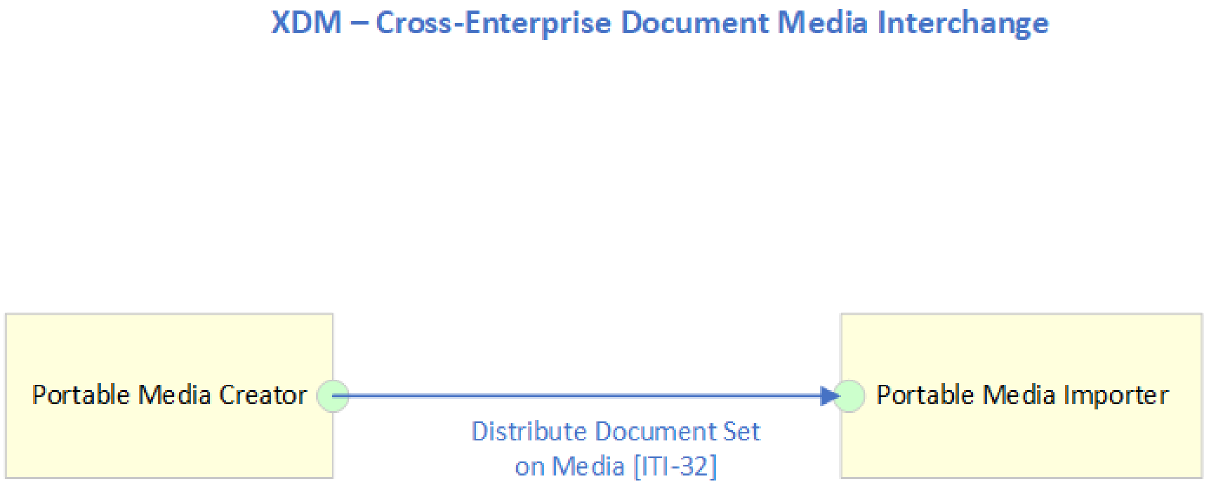
XDM – Cross-Enterprise Document Media Interchange
Portable Media Creator
Distribute Document Set on Media [ITI-32]
Portable Media Importer

The table below lists the transactions for each actor directly involved in the XDM profile. To claim compliance with XDM, an actor shall support all required transactions (labeled "R").

Actor	Transaction	Optionality
Portable Media Creator	Distribute Document Set on Media [ITI-32]	R
Portable Media Importer	Distribute Document Set on Media [ITI-32]	R

Transactions

- Distribute Document Set on Media [ITI-32] - the Portable Media Creator sends information to media reading actors by means of Interchange Media (email, USB drive, CD/DVD ROM) where it stores the information.

Sequence Diagram

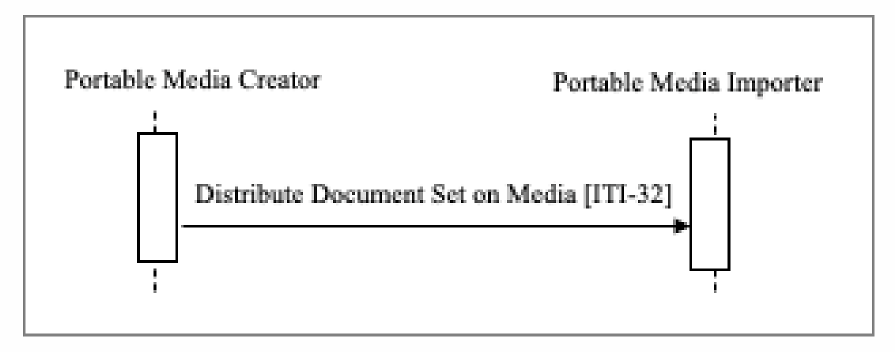

1.5 Patient Mediated Access (PMA) Profiles

The pan-Canadian Patient Mediated Access (PMA) is a set of pan-Canadian specifications aiming to create a modern healthcare ecosystem and promote the building blocks in which patients have the control and authority to securely manage and share their health information with authorized healthcare providers and caregivers.

1.5.1 Shareable Health Links (CA:SHL)

Background

The CA:SHL specification is a Canadian Interoperability Specifications (indicated by the CA: prefix) based on the HL7 Health Links Specification, currently in the process of being contributed to IHE (Integrating the Healthcare Enterprise), with the aim of becoming an internationally adopted Integration Profile. As part of this transition, the names of the profiles, actors, and transactions are subject to change and align with international standards and conventions.

Note that Shareable Health Link is based on HL7 SMART Health Links and often can be used interchangeably, Shareable Health Links offers an Integration Profile around SMART Health Link, with the goal to provide implementers with a set of building blocks using the Actor and Transaction paradigm, and a consistent way to conformance test their implementations.

Security Considerations

When creating, sharing, and consuming clinical data using CA:SHL, it is crucial to also ensure the privacy, authenticity, and integrity of the information through encryption and digital signatures.

- Specific details regarding the privacy, authenticity, and integrity of the information are not within the scope of this specification. However, implementations that do not employ high-level security measures, such as digital signatures and robust encryption, are not recommended and are strongly discouraged.
- CA:SHL is the subject of international collaboration at IHE International, with the IHE ITI committee approving the initiative to evolve SHLinks and advance it into Verifiable Health Links, by adding a structured and robust security approach in future iterations.
- While trust framework considerations are critical, they fall outside the scope of this specification. Implementers should ensure that any exchange of SHLinks is governed by a robust trust framework that ensure security and provenance.

- CA:SHL is a building block that is meant to be used together with added security measures, otherwise it is not suitable for exchange in environments where security and provenance cannot be reliably established by other means.

Standard Version

The CA:SHL Profile is based on:

HL7 Health Links Specification: version 1.0.0-ballot (2024-08-08), based on FHIR v4.0.1: R4. This is part of the proposed STU1 version of the SMART Health Cards and Links FHIR IG for ballot in September 2024.

Acronyms

CA:SHL – Canadian Shareable Health Links (CA:SHL) Interoperability Specifications

SHL, SHLink – Shareable Health Link, containing access point to the patient's health data (often encoded in a QR code format).

SHLink Clinical Data - Clinical data associated with the SHLink

Overview

Shared Health Links (CA:SHL) is a Canadian Interoperability Specifications based on the HL7 Health Links Specification, that defines building blocks enabling patients to generate shareable links that encode their health data such as patient summaries or immunization records. These shareable links can often be downloaded onto their devices and converted in a QR code format, facilitating patient-mediated data sharing and interoperability within the healthcare ecosystem.

The following are some examples of the benefits of SHLink:

- Enables patients to create, share, and manage their health data using standardized operations based on HL7 standards.
- Facilitates convenient access to health information through easily shareable links, such as QR codes.
- Enhances patient control over their health data, allowing them to decide who can access their information.
- Promotes seamless coordination and collaboration among healthcare providers by providing timely and accurate health information.
- Supports improved continuity of care through efficient data sharing during patient transitions across different care settings.

Actors and Transactions

The following diagram provides an overview of the CA:SHL Actors, Transactions and their interactions that are currently in scope.

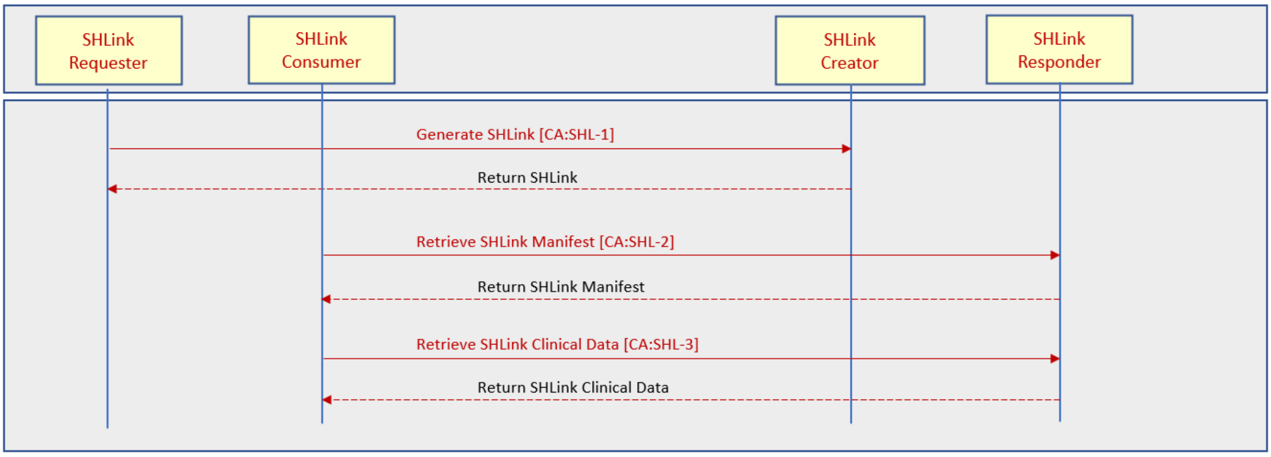
SHLink Requester
SHLink Consumer
SHLink Creator
SHLink Responder
Generate SHLink [CA:SHL-1]
Return SHLink
Retrieve SHLink Manifest [CA:SHL-2]
Return SHLink Manifest
Retrieve SHLink Clinical Data [CA:SHL-3]
Return SHLink Clinical Data

The table below lists the transactions for each actor directly involved in CA:SHL. To claim compliance with CA:SHL, an actor shall support all required transactions (labeled “R”).

Actor	**Transaction**	**Optionality**
SHLink Requester	Generate SHLink [CA:SHL-1]	R
SHLink Creator	Generate SHLink [CA:SHL-1]	R
SHLink Consumer	Retrieve SHLink Manifest [CA:SHL-2]	R
	Retrieve SHLink Clinical Data [CA:SHL-3]	R
SHLink Responder	Retrieve SHLink Manifest [CA:SHL-2]	R
	Retrieve SHLink Clinical Data [CA:SHL-3]	R

Transaction Details

- Generate SHLink [CA:SHL-1] – This transaction sends a Generate SHLink Request from an SHLink Requester to an SHLink Creator, using an HTTP POST operation. The SHLink Creator returns an SHLink that contains access to the manifest file for accessing the clinical data that was requested.
- Retrieve SHLink Manifest [CA:SHL-2] – This transaction sends a Retrieve SHLink Manifest Request from an SHLink Consumer to an SHLink Responder. The SHLink Responder returns the Manifest file that contains the location of the clinical data that was requested.
- Retrieve SHLink Clinical Data [CA:SHL-3] – This transaction sends a Retrieve SHLink Clinical Data Request from an SHLink Consumer to an SHLink Responder. The SHLink Responder returns the clinical data that was requested with applied encryption. The SHLink Consumer decrypts the data using the key provided in the SHLink.

CA:SHL - Canadian Implementation Guidance

Scope

The initial phase of CA:SHL focuses on a subset of the standards, that can be used to build a Minimum Viable Product.

Use Cases

There are two basic Use Cases identified that are currently in scope of this specification:

1. Generate SHLink
2. Consume SHLink and access clinical data

These two use cases are described in detail in the pan-Canadian Patient Summary (PS-CA) specifications, where the patient presents their SHLink to a Health Care Provider, the clinical data that is shared being their Patient Summary.

Note that currently the information included into the CA:SHL is minimal, as the CA:SHL profile is in the process of being contributed and evolved as part of an international collaboration at IHE International,

Implementation Details

Since the CA:SHL specification is currently minimal and in its early stages, it will evolve into a more comprehensive specification through collaboration with IHE and potentially other international partners.

The CA:SHL Canadian Implementation Guidance describes the recommendations below for early-stage implementations.

SHLink payload	• Makes use of a manifest file (`url` element) • Includes decryption key for processing clinical data returned in manifest (`key` element) • Passcode is mandatory (`P` flag) • Bypassing the manifest is not supported (`U` flag) • Contains an expiration time (`exp` element)
Manifest file	• Contains `.files.contentType: application/fhir+json` • Contains `.files.location` • `.files.embedded` content is not included
Clinical data	• Encryption is applied using symmetrical JSON Web Encryption (JOSE JWE)

SHLink format	• SHLink should be encoded into a QR code format (on the client side) • SHLink should contain a viewer URL that ends with #
Under exploration	• Standard endpoints for generating and other operations related to SHLinks • Details around specifying different types of data to be shared (e.g. Patient Summary, Immunization, etc.)

1.6 Patient Identity Profiles

The following Patient Identity Profiles are included:

- Patient Master Identity Registry (PMIR)
- Patient Identifier Cross-reference for Mobile (PIXm)
- Patient Demographics Query for Mobile (PDQm)

1.6.1 Patient Master Identity Registry (PMIR)

Overview

The Patient Master Identity Registry (PMIR) Profile supports creating, updating, and deprecating patient identity information about a subject of care, as well as subscribing to changes, using HL7 FHIR resources and RESTful transactions. This profile includes the Patient Identifier Cross-reference for Mobile (PIXm) and Patient Demographics Query for Mobile (PDQm) profiles. The "patient master identity" is the dominant patient identity managed centrally among many participating organizations (a.k.a., "Golden Patient Identity").

Beyond the basic create, retrieve, update, and delete transaction set, this profile addresses important patient safety issues related to cases where there are two or more patient master identities that have been established for the same person, thus it is not clear which identity is the "true" one. There is also a risk that health data (possibly conflicting) may be associated with each identity – and these disparate data, together, may need to be reconciled before a fully and accurate "health picture" can be developed for this person. These situations represent patient safety risks. This profile addresses how these multiple patient master identities can be merged into a single patient master identity, and how this merge flows down to data custodians so that they take appropriate actions. It is outside the scope of this profile to define how references to the deprecated patient master identity from other data should be handled.

Actors and Transactions

The following diagram provides an overview of the PMIR profile Actors, Transactions and their interactions.

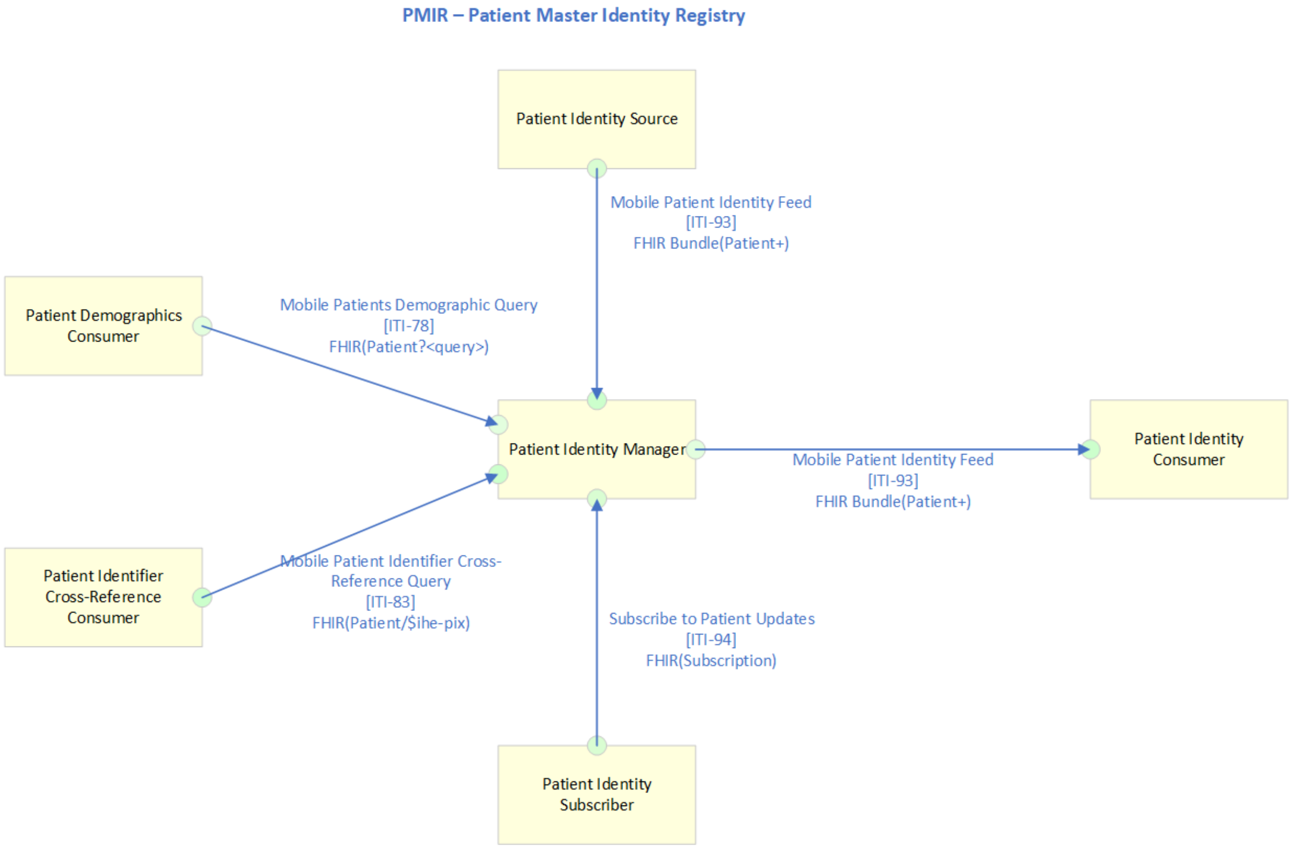
PMIR – Patient Master Identity Registry
Patient Identity Source
Mobile Patient Identity Feed
[ITI-93]
FHIR Bundle(Patient+)
Patient Demographics Consumer
Mobile Patients Demographic Query
[ITI-78]
FHIR(Patient?<query>)
Patient Identity Manager
Mobile Patient Identity Feed
[ITI-93]
FHIR Bundle(Patient+)
Patient Identity Consumer
Patient Identifier Cross-Reference Consumer
Mobile Patient Identifier Cross-Reference Query
[ITI-83]
FHIR(Patient/$ihe-pix)
Subscribe to Patient Updates
[ITI-94]
FHIR(Subscription)
Patient Identity Subscriber

The table below lists the transactions for each actor directly involved in the PMIR profile. To claim compliance with PMIR, an actor shall support all required transactions (labeled “R”).

Actor	Transaction	Optionality
Patient Identity Source	Mobile Patient Identity Feed [ITI-93]	R
Patient Identity Consumer	Mobile Patient Identity Feed [ITI-93]	R
Patient Identity Manager	Mobile Patient Identity Feed [ITI-93] Mobile Patient Identifier Cross-Reference Query [ITI-83] Mobile Patient Demographic Query [ITI-78] Subscribe to Patient Updates [ITI-94]	R R R R
Patient Demographics Consumer	Mobile Patient Demographic Query [ITI-78]	R
Patient Identity Cross-Reference Consumer	Mobile Patient Identifier Cross-Reference Query [ITI-83]	R

Actor	Transaction	Optionality
Patient Identity Subscriber	Subscribe to Patient Updates [ITI-94]	R

Transactions

- Mobile Patient Demographic Query [ITI-78] - Requests a list of patients matching the supplied set of demographics criteria (example: ID or Name) from the Patient Demographics Supplier. The Patient Demographics Consumer populates its attributes with demographic information received from the Patient Demographics Supplier. Uses a FHIR Patient Query: GET [base]/Patient?<query>
- Mobile Patient Identifier Cross-Reference Query [ITI-83] - Used by the Patient Identity Cross-Reference Consumer actor to solicit information about patients whose Patient Identifiers cross-match with Patient Identifiers provided in the query parameters of the request message. The request is received by the Patient Identifier Cross-Reference Manager. The Patient Identifier Cross-Reference Manager processes the request and returns a response in the form of zero or more Patient Identifiers for the matching patient. Uses a FHIR Patient $ihe-pix operation: GET [base]/Patient/$ihe-pix?<query>
- Mobile Patient Identity Feed [ITI-93] – Used to send a FHIR Bundle of new and updated Patient Resources. POSTs a FHIR Bundle containing at least one Patient resource.
- Subscribe to Patient Updates [ITI-94] - allows a Patient Identity Subscriber to subscribe to a Mobile Patient Resource Feed [ITI-93] depending on the requested criteria. POSTs a FHIR Subscription resource.

Sequence Diagrams

PMIR Subscriber and Registry

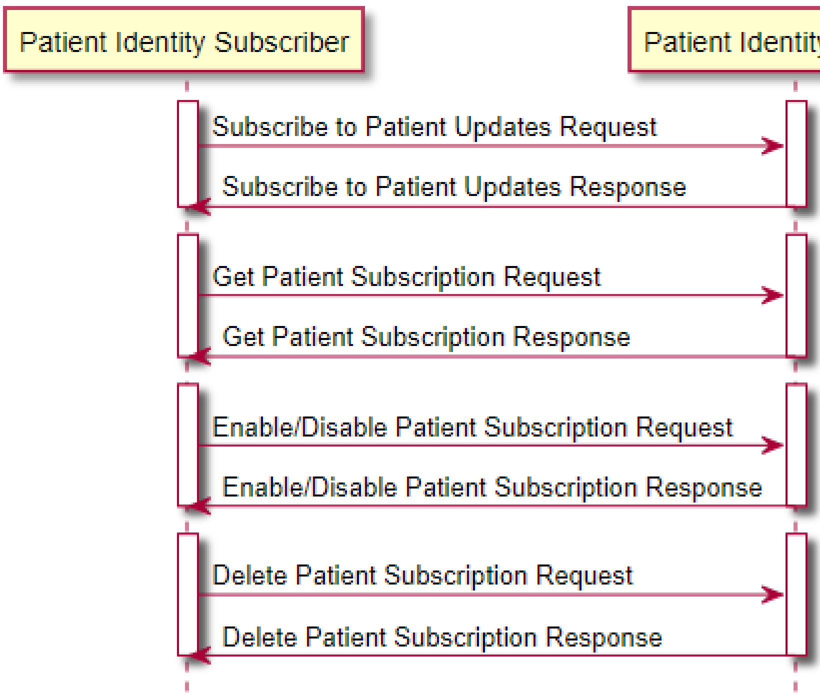
Patient Identity Subscriber
Patient Identity Registry
Subscribe to Patient Updates Request
Subscribe to Patient Updates Response
Get Patient Subscription Request
Get Patient Subscription Response
Enable/Disable Patient Subscription Request
Enable/Disable Patient Subscription Response
Delete Patient Subscription Request
Delete Patient Subscription Response

PMIR Supplier and Consumer

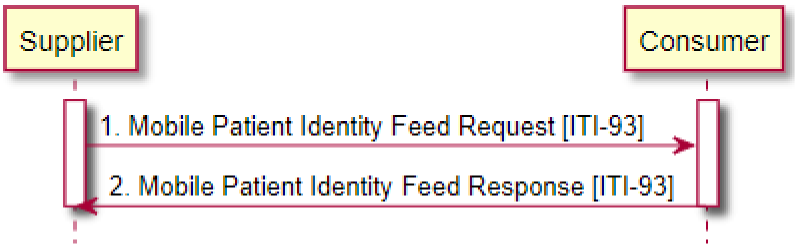
Supplier
Consumer
1. Mobile Patient Identity Feed Request [ITI-93]
2. Mobile Patient Identity Feed Response [ITI-93]

PMIR Create Patient Identity

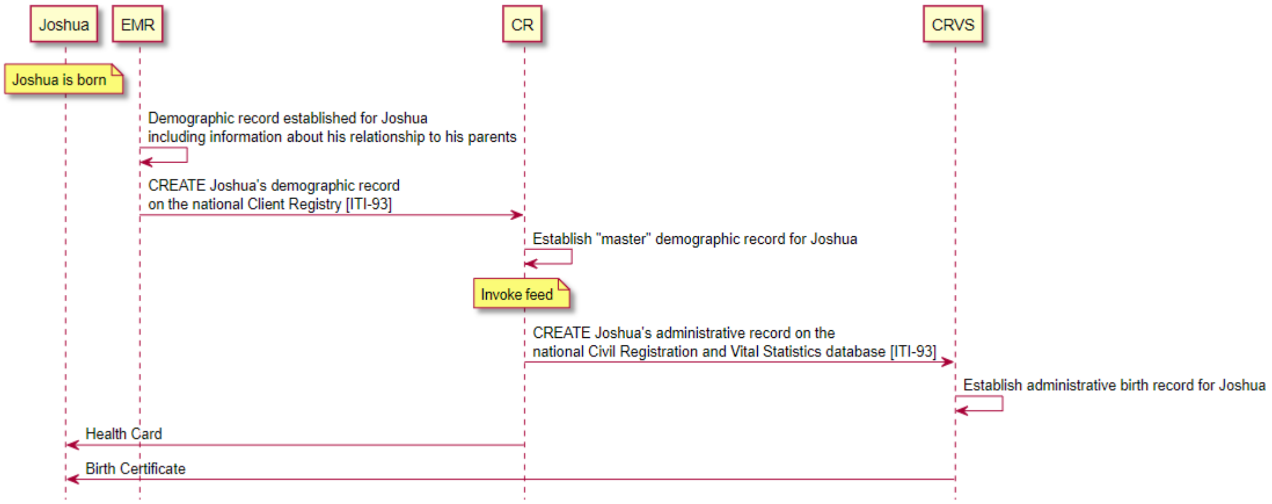
Joshua
EMR
CR
CRVS
Joshua is born
Demographic record established for Joshua
including information about his relationship to his parents
CREATE Joshua's demographic record
on the national Client Registry [ITI-93]
Establish "master" demographic record for Joshua
Invoke feed
CREATE Joshua's administrative record on the
national Civil Registration and Vital Statistics database [ITI-93]
Establish administrative birth record for Joshua
Health Card
Birth Certificate

PMIR Update Patient Identity

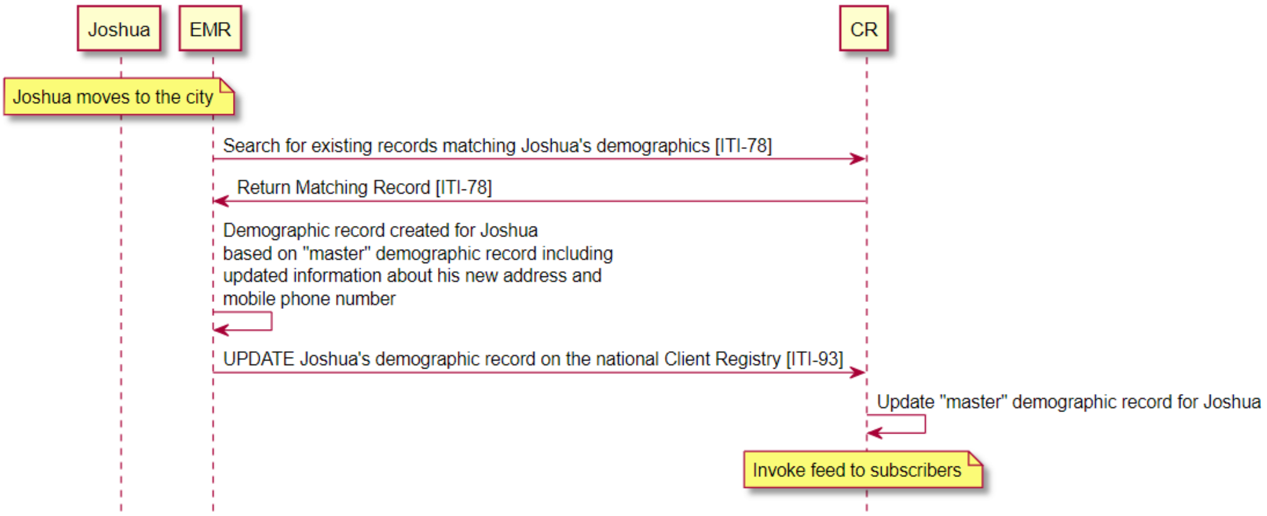
Joshua
EMR
CR
Joshua moves to the city
Search for existing records matching Joshua's demographics [ITI-78]
Return Matching Record [ITI-78]
Demographic record created for Joshua based on "master" demographic record including updated information about his new address and mobile phone number
UPDATE Joshua's demographic record on the national Client Registry [ITI-93]
Update "master" demographic record for Joshua
Invoke feed to subscribers

PMIR Merge Patient Identity

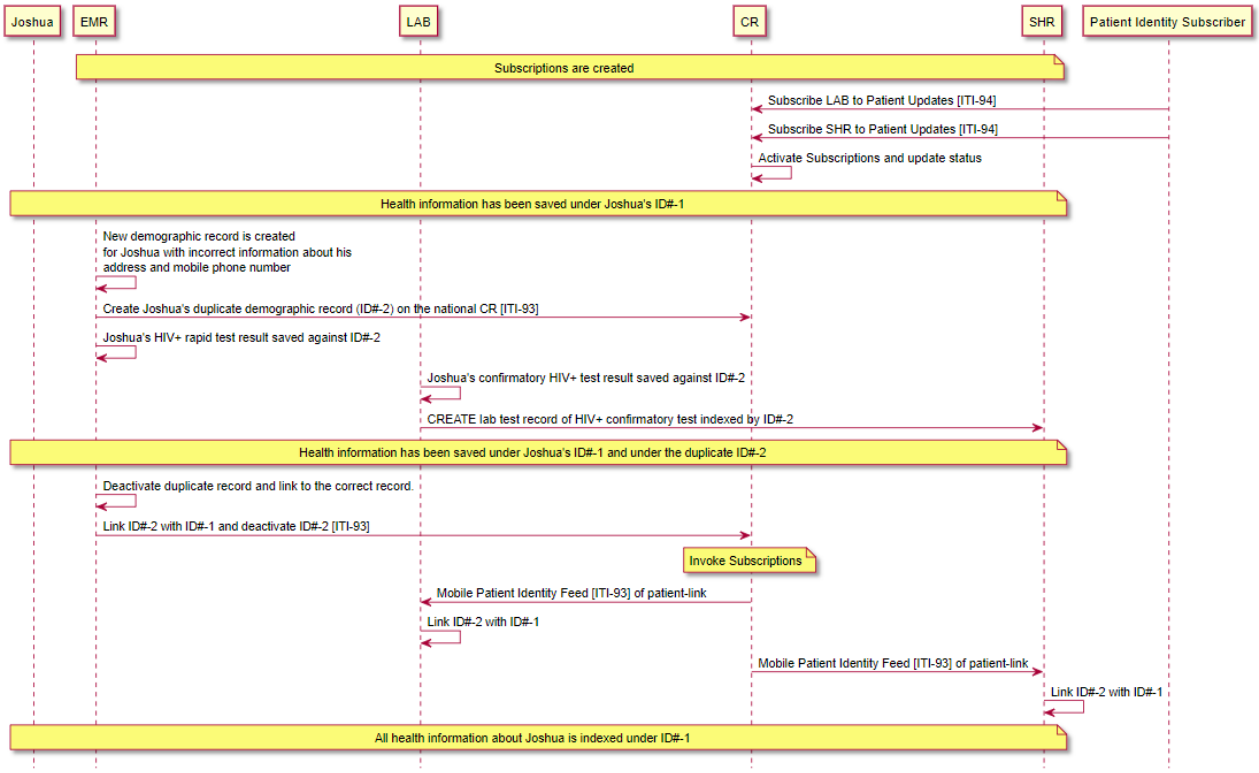
Joshua
EMR
LAB
CR
SHR
Patient Identity Subscriber
Subscriptions are created
Subscribe LAB to Patient Updates [ITI-94]
Subscribe SHR to Patient Updates [ITI-94]
Activate Subscriptions and update status
Health information has been saved under Joshua's ID#-1
New demographic record is created for Joshua with incorrect information about his address and mobile phone number
Create Joshua's duplicate demographic record (ID#-2) on the national CR [ITI-93]
Joshua's HIV+ rapid test result saved against ID#-2
Joshua's confirmatory HIV+ test result saved against ID#-2
CREATE lab test record of HIV+ confirmatory test indexed by ID#-2
Health information has been saved under Joshua's ID#-1 and under the duplicate ID#-2
Deactivate duplicate record and link to the correct record.
Link ID#-2 with ID#-1 and deactivate ID#-2 [ITI-93]
Invoke Subscriptions
Mobile Patient Identity Feed [ITI-93] of patient-link
Link ID#-2 with ID#-1
Mobile Patient Identity Feed [ITI-93] of patient-link
Link ID#-2 with ID#-1
All health information about Joshua is indexed under ID#-1

1.6.2 Patient Identifier Cross-reference for Mobile (PIXm)

Overview

The Patient Identifier Cross-Reference for Mobile (PIXm) Integration Profile provides a transaction for mobile and lightweight browser-based applications to query a Patient Identifier Cross-Reference Manager for a list of patient identifiers based on the patient identifier in a different domain and retrieve a patient's cross-domain identifiers information into the application.

PIXm is intended to be used by lightweight applications and mobile devices present in healthcare enterprises of a broad range of sizes (hospital, a clinic, a physician office, etc.). It supports the cross-reference query of patient identifiers from multiple Patient Identifier Domains by providing the ability to access the list(s) of cross-referenced patient identifiers via a query/response.

Actors and Transactions

The following diagram provides an overview of the PIXm profile Actors, Transactions and their interactions.

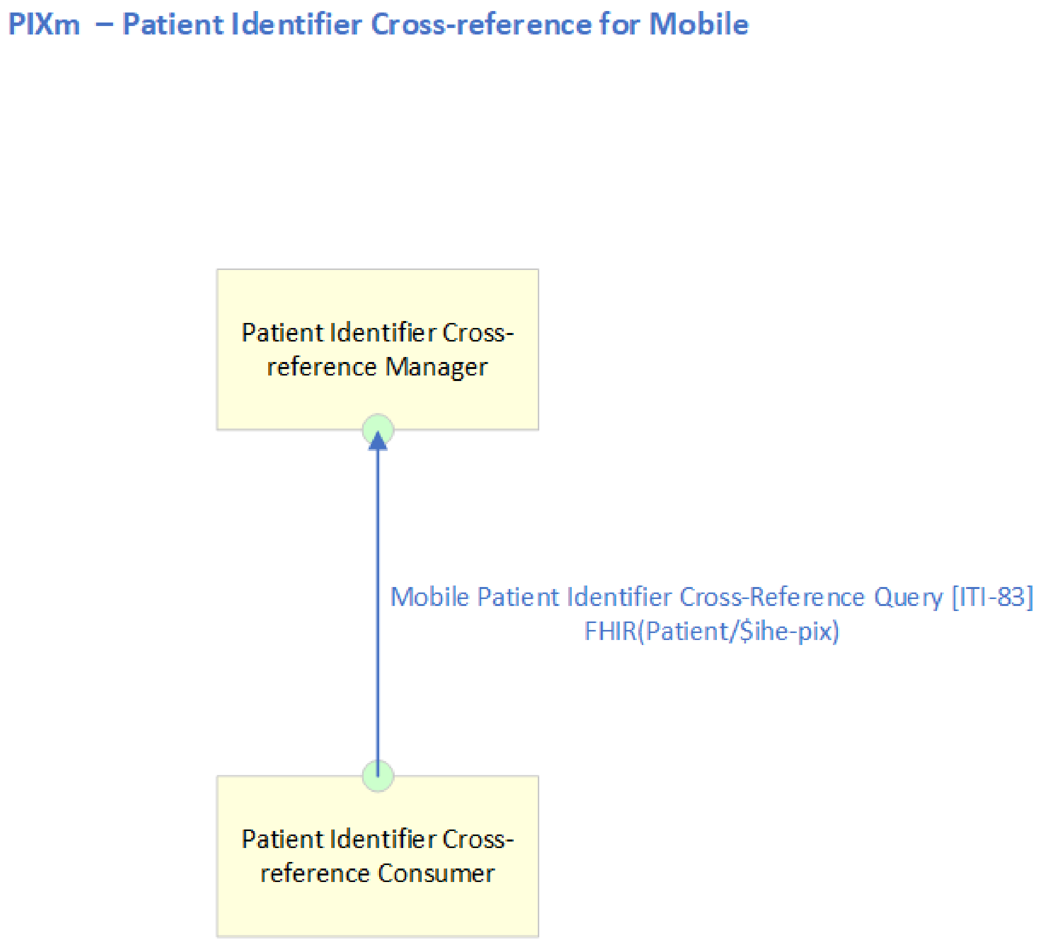
PIXm – Patient Identifier Cross-reference for Mobile
Patient Identifier Cross-reference Manager
Mobile Patient Identifier Cross-Reference Query [ITI-83]
FHIR(Patient/$ihe-pix)
Patient Identifier Cross-reference Consumer

The table below lists the transactions for each actor directly involved in the PIXm profile. To claim compliance with PIXm, an actor shall support all required transactions (labeled "R").

Actor	Transaction	Optionality
Patient Identity Manager	Mobile Patient Identifier Cross-Reference Query [ITI-83]	R
Patient Identity Cross-Reference Consumer	Mobile Patient Identifier Cross-Reference Query [ITI-83]	R

Transactions

- Mobile Patient Identifier Cross-Reference Query [ITI-83] - Used by the Patient Identity Cross-Reference Consumer actor to solicit information about patients whose Patient Identifiers cross-match with Patient Identifiers provided in the query parameters of the request message. The request is received by the Patient Identifier Cross-Reference Manager. The Patient Identifier Cross-Reference Manager processes the request and returns a response in the form of zero or more Patient Identifiers for the matching patient. Uses a FHIR Patient $ihe-pix operation: GET [base]/Patient/$ihe-pix?<query>

Sequence Diagram

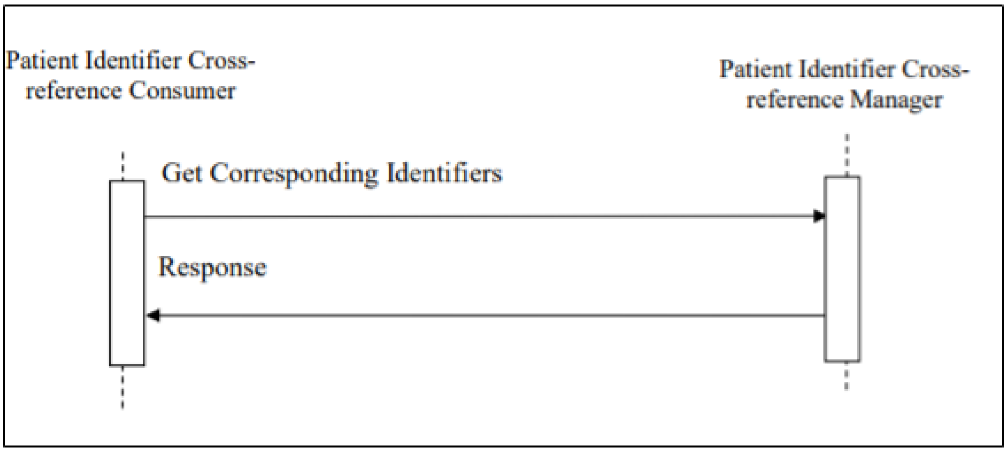
Patient Identifier Cross-reference Consumer
Patient Identifier Cross-reference Manager
Get Corresponding Identifiers
Response

1.6.3 Patient Demographics Query for Mobile (PDQm)

Overview

The Patient Demographics Query for Mobile (PDQm) Profile provides a transaction for mobile and lightweight browser-based applications to query a patient demographics supplier for a list of patients based on user-defined search criteria and retrieve a patient’s demographic information.

Using these patterns, the PDQm Profile exposes the functionality of a patient demographics supplier to mobile applications and lightweight browser applications. The following list provides a few examples of how PDQm might be leveraged by implementers:

- A health portal securely exposing demographics data to browser based plugins
- Medical devices which need to access patient demographic information
- Mobile devices used by physicians (example bedside eCharts) which need to establish patient context by scanning a bracelet
- Web based EHR/EMR applications which wish to provide dynamic updates of patient demographic information such as a non-postback search, additional demographic detail, etc.
- Any low resource application which exposes patient demographic search functionality
- Any application using the MHD Profile to access documents may use PDQm to find an appropriate patient identifier

Actors and Transactions

The following diagram provides an overview of the PDQm profile Actors, Transactions and their interactions.

PDQm – Patient Demographics Query for Mobile

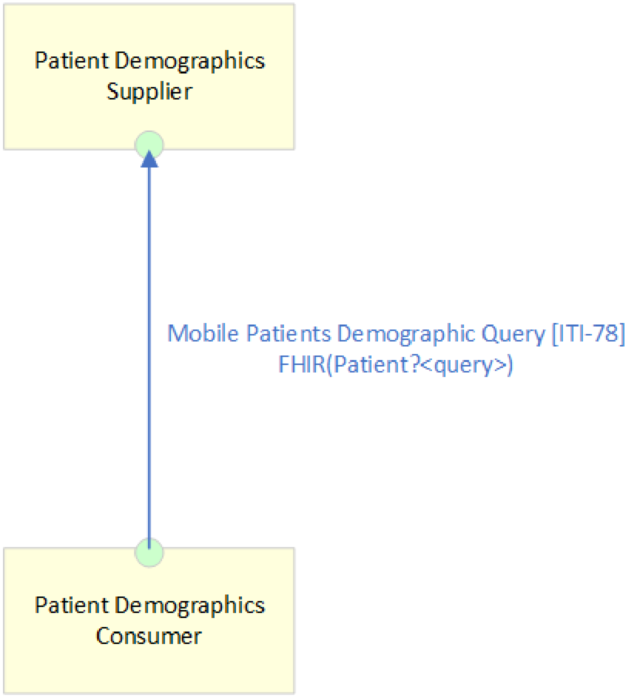
Patient Demographics Supplier
Mobile Patients Demographic Query [ITI-78]
FHIR(Patient?<query>)
Patient Demographics Consumer

The table below lists the transactions for each actor directly involved in the PDQm profile. To claim compliance with PDQm, an actor shall support all required transactions (labeled “R”).

Actor	Transaction	Optionality
Patient Identity Manager	Mobile Patient Demographic Query [ITI-78]	R
Patient Demographics Consumer	Mobile Patient Demographic Query [ITI-78]	R

Transactions

- Mobile Patient Demographic Query [ITI-78] - Requests a list of patients matching the supplied set of demographics criteria (example: ID or Name) from the Patient Demographics Supplier. The Patient Demographics Consumer populates its attributes with demographic information received from the Patient Demographics Supplier. Uses a FHIR Patient Query: GET [base]/Patient?<query>

Sequence Diagram

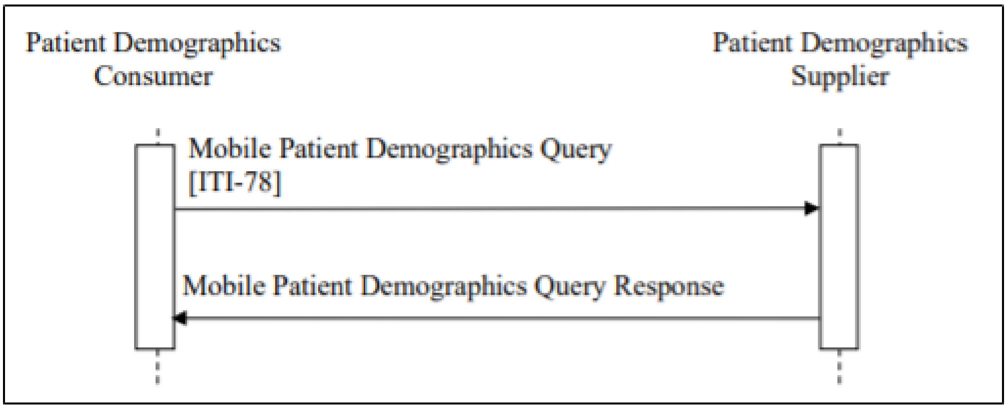
Patient Demographics Consumer
Patient Demographics Supplier
Mobile Patient Demographics Query [ITI-78]
Mobile Patient Demographics Query Response