



**HL7 White Paper: Retrieval of Immunization Data via
Bulk FHIR, Edition 1- US Realm**

January 2025

HL7 Informative Ballot

**Sponsored by:
Public Health Work Group**

Copyright © 2024-2025 Health Level Seven International ® ALL RIGHTS RESERVED. The reproduction of this material in any form is strictly forbidden without the written permission of the publisher. HL7 and Health Level Seven are registered trademarks of Health Level Seven International. Reg. U.S. Pat & TM Off.

Use of this material is governed by HL7's [IP Compliance Policy](#).

IMPORTANT NOTES:

HL7 licenses its standards and select IP free of charge. **If you did not acquire a free license from HL7 for this document**, you are not authorized to access or make any use of it. To obtain a free license, please visit <http://www.HL7.org/implement/standards/index.cfm>.

If you are the individual that obtained the license for this HL7 Standard, specification or other freely licensed work (in each and every instance "Specified Material"), the following describes the permitted uses of the Material.

A. HL7 INDIVIDUAL, STUDENT AND HEALTH PROFESSIONAL MEMBERS, who register and agree to the terms of HL7's license, are authorized, without additional charge, to read, and to use Specified Material to develop and sell products and services that implement, but do not directly incorporate, the Specified Material in whole or in part without paying license fees to HL7.

INDIVIDUAL, STUDENT AND HEALTH PROFESSIONAL MEMBERS wishing to incorporate additional items of Special Material in whole or part, into products and services, or to enjoy additional authorizations granted to HL7 ORGANIZATIONAL MEMBERS as noted below, must become ORGANIZATIONAL MEMBERS of HL7.

B. HL7 ORGANIZATION MEMBERS, who register and agree to the terms of HL7's License, are authorized, without additional charge, on a perpetual (except as provided for in the full license terms governing the Material), non-exclusive and worldwide basis, the right to (a) download, copy (for internal purposes only) and share this Material with your employees and consultants for study purposes, and (b) utilize the Material for the purpose of developing, making, having made, using, marketing, importing, offering to sell or license, and selling or licensing, and to otherwise distribute, Compliant Products, in all cases subject to the conditions set forth in this Agreement and any relevant patent and other intellectual property rights of third parties (which may include members of HL7). No other license, sublicense, or other rights of any kind are granted under this Agreement.

C. NON-MEMBERS, who register and agree to the terms of HL7's IP policy for Specified Material, are authorized, without additional charge, to read and use the Specified Material for evaluating whether to implement, or in implementing, the Specified Material, and to use Specified Material to develop and sell products and services that implement, but do not directly incorporate, the Specified Material in whole or in part.

NON-MEMBERS wishing to incorporate additional items of Specified Material in whole or part, into products and services, or to enjoy the additional authorizations granted to HL7 ORGANIZATIONAL MEMBERS, as noted above, must become ORGANIZATIONAL MEMBERS of HL7.

Please see <http://www.HL7.org/legal/ippolicy.cfm> for the full license terms governing the Material.

Ownership. Licensee agrees and acknowledges that **HL7 owns** all right, title, and interest, in and to the Materials. Licensee shall **take no action contrary to, or inconsistent with**, the foregoing.

Licensee agrees and acknowledges that HL7 may not own all right, title, and interest, in and to the Materials and that the Materials may contain and/or reference intellectual property owned by third parties ("Third Party IP"). Acceptance of these License Terms does not grant Licensee any rights with respect to Third Party IP. Licensee alone is responsible for identifying and obtaining any necessary licenses or authorizations to utilize Third Party IP in connection with the Materials or otherwise. Any actions, claims or suits brought by a third party resulting from a breach of any Third Party IP right by the Licensee remains the Licensee's liability.

Following is a non-exhaustive list of third-party terminologies that may require a separate license:

Terminology	Owner/Contact
Current Procedures Terminology (CPT) code set	American Medical Association https://www.ama-assn.org/practice-management/cpt-licensing
SNOMED CT®	SNOMED CT® International; http://www.snomed.org/snomed-ct/get-snomed-ct or info@ihtsdo.org
Logical Observation Identifiers Names & Codes (LOINC®)	Regenstrief Institute
International Classification of Diseases (ICD) codes	World Health Organization (WHO)
NUCC Health Care Provider Taxonomy code set	American Medical Association. Please see www.nucc.org . AMA licensing contact: 312-464-5022 (AMA IP services)

Helios Bulk Data for Immunization Guidance Document

The purpose of this document is to identify and describe the use of the HL7 FHIR Bulk Data Access approach to sharing data on populations of individuals. The Bulk Data approach can be used to ensure that users of Electronic Health Records (EHR) systems and other electronic systems can benefit from the carefully collated and curated data entrusted to public health programs. By implementing this strategy, authorized users can efficiently retrieve data for defined set of individuals through the use of a standardized and reusable approach and allows public health systems to generate and transmit the necessary information swiftly and with a minimum amount of time and effort expended on responding to data requests using proprietary and bespoke formats.

Helios Bulk Data for Immunization Guidance Document.....	3
Introduction	5
What is the purpose of this document?	5
What is the FHIR Bulk Data Approach?.....	5
Which Public Health Use Cases Could Utilize Bulk Data?.....	5
Who is Involved?	6
What are the Steps in Accessing Bulk Data?	7
What are the Benefits of Bulk Data Access?	7
Policy Considerations.....	8
Technical Concepts	10
Bulk Data Access FHIR IG Overview	10
Security Expectations	13
Technology Implementation	13
Group Types.....	16
Maintaining Cohorts	18
Data Payload	20
Trading Partner Expectations	20
Foundational Technical Requirements.....	21
Foundational Operational Requirements	21
Mechanisms for Patient Identification and Group Creation	22
V2-Mediated Matching (Automatic Group Creation).....	22
V2-Mediated Matching (Dynamic Group Creation)	26
FHIR-Mediated Matching (Individual).....	29
FHIR-Mediated Matching (Bulk).....	32
Flat File Mediated Matching	35
Reuse of Existing Groupings	38
Online Interface.....	40
Subscription to Patient with Bulk	42
Mechanism for Bulk Query	45
Summary	45

Introduction

What is the purpose of this document?

This document is the product of the [Helios FHIR Accelerator](#) for Public Health [Bulk Data Priority Area](#). This group of volunteers has explored the potential uses of Bulk Data exchange to understand the approaches that trading partners may take and identify considerations that must be taken into account by implementers. Note that while this document has been developed for the US Realm, this discussion within may be broadly applicable to other locales as well.

So that both public health programs and authorized users can better understand the methodology and implications of the Bulk Data Access approach, in this document we discuss:

- Policy considerations and implications of implementing Bulk Data Access using FHIR,
- Technical concepts and options for developing a supporting FHIR infrastructure,
- Trading partner expectations,
- Options for identifying and grouping individuals,
- And approaches for executing the bulk query itself.

What is the FHIR Bulk Data Approach?

The [FHIR Bulk Data Access Implementation Guide](#) describes a generic way for trading partners to exchange large volumes of information on populations of individuals. This approach may be applied to a wide variety of public health data sets and a number of different Authorized User types including healthcare providers, payors, and community groups such as schools or community-based program providers. Regardless of the nature of the individual actors involved, the application of the FHIR Bulk Data approach provides a consistent framework for data exchange to minimize the use of bespoke formats. We will provide a more detailed description of this approach in the [Technical Concepts](#) section of this document

Which Public Health Use Cases Could Utilize Bulk Data?

General Use Cases

As part of their overall mission and objectives, many public health programs collect, collate and curate valuable data sets. This information is of use not only to public health but also to others who directly interact with both small and large populations of individuals to provide healthcare, services & support and other interventions. By offering authorized users improved access to these data sets, public health can promote improved health and quality of life to members of their community.

The adoption of FHIR-mediated bulk data exchange by public health will promote expanded data sharing while minimizing time and effort on the part of both authorized users and public health. A wide variety of public health programs are stewards of data which may be of use to

healthcare providers, payers, service providers, community programs, other public health programs and more. A few example use cases discussed within the Helios Bulk Data Priority Area include:

- Vital Records birth certification data, which may be directly useful to other public health programs tasked with monitoring newborns and providing support and services.
- Vital Records death data, which would allow service providers such as healthcare organizations and payers to update their data on individuals who have expired, reducing costs associated with unnecessary resource allocations and the emotional toll on survivors due to inadvertent communication.
- Exchange of data between state and local public health programs for geographically located populations.
- Public health collated data which may support care coordination with a Medicaid agency.

However, the primary use case explored within the Bulk Data Priority Area is the sharing of immunization-related data collected by jurisdictional immunization programs within an Immunization Information System (IIS). This use case is explored more thoroughly below.

Immunization-Specific Use Cases

As discussed in the [Immunization Integration Program's Bulk Query Toolkit](#), the following immunization-specific use cases may benefit from employing FHIR bulk data exchange rather than individual queries:

- Obtaining immunization records for all patients with clinic appointments the following day
- *Limiting* reminder or recall communication to patients who are due or overdue for a vaccine dose
- Evaluation of clinical care quality through performance measures (e.g., Healthcare Effectiveness Data and Information Set - HEDIS)
- Supporting public health agency disease case investigations by identifying exposure of unvaccinated persons
- Evaluating the vaccination status of large groups of known people (i.e., coverage rates for health plans, schools, or emergency first responders)
- Providing consumers with proof of vaccination for employment, travel, education, or large events
- Evaluating vaccine product effectiveness after FDA and ACIP approval

Who is Involved?

Four major partner roles are anticipated to be involved in the development and implementation of Bulk Data exchange for immunizations. In later sections of this document, we will discuss specific expectations and activities required of each of these roles during the development and implementation process. These major roles are:

- Public health program

- Health Information Technology (HIT) vendors for public health systems
- Authorized User organization
- Authorized User HIT vendor

For the purposes of this document, we will use a jurisdictional immunization program as an example, but similar approaches and considerations can be applied to other public health programs. In this example, the HIT vendor for the immunization program would be the Immunization Information System (IIS) vendor.

For the immunization use case, many different types of organizations might be considered an Authorized User. These include:

- Healthcare providers (e.g., hospitals and clinics)
- Payers
- Community-based organizations (e.g., schools, community programs)
- Other public health programs
 - Within the jurisdiction
 - Within other jurisdictions such as a neighboring state or local agency
- Others as determined by the public health program sharing the data

The needs and technical capabilities of different types of Authorized Users (and even within a given class of Authorized User) are expected to vary greatly. As such, not all approaches discussed in this document will be practical for all types of Authorized Users. Implementers are expected to fully understand the needs and capabilities of a trading partner and design an approach which can be practically implemented on both sides of the exchange.

What are the Steps in Accessing Bulk Data?

The application of the Bulk Data Access approach by an Authorized User is expected to consist of multiple steps including:

- Patient identification between systems for a significant numbers of individuals
- Creation and maintenance of cohorts of individuals
- Use of a FHIR operation for bulk data exchange

Note that only the last of these steps is covered in the [FHIR Bulk Data Access Implementation Guide](#) itself which assumes the existence of an agreed upon cohort of individuals known by both systems. In this guidance document, we will also consider approaches aimed at the identification of individuals at scale and the creation and maintenance of cohorts of individuals to meet these necessary prerequisites to performing a bulk query.

What are the Benefits of Bulk Data Access?

The use of FHIR-based Bulk Data queries brings benefits to both the data source as well as the Authorized Users of that data who request access. For example, in the Immunization space

there already exist standard methods for querying and receiving data via HL7 v2 messaging, but this approach only supports patient queries on a one-by-one basis. The v2 standards are optimized for accessing data on an individual basis, and accessing data on larger numbers of individuals can put large operational burdens on the IIS which must respond to these queries synchronously (immediately). Furthermore, during the COVID pandemic, immunization programs were additionally faced with requests to pull data in bulk on large numbers of individuals. Even when the IIS offered data at scale in a standard format, these requests were typically through bespoke flat file-based mechanisms and each request required individual planning, coordination and often custom development, creating a bottleneck. When balanced against other priorities, public health staff often did not have the necessary capacity to respond to all requests for data. However, FHIR Bulk offers the opportunity to respond to large data access through a standardized and replicable approach in the same way HL7 v2 has standardized clinical access for individuals. This will allow for consistent and secure access to data for Authorized Users, using standards-based paradigms. This will help the immunization programs operate less as a silo of data and become more of a data platform.

Furthermore, many Authorized Users such as healthcare providers or payers operate across different jurisdictions and request data from multiple data sources. In the past, variation by jurisdiction has been a significant burden on Authorized Users and their HIT vendors as each jurisdiction would often require customized solutions for exchanging data. By uniformly adopting FHIR-based Bulk Data exchange, the national landscape will become more harmonized for developers and implementers.

As another benefit, FHIR provides a common data structure to represent immunization data. By leveraging a Bulk Data query approach, immunization data will have a similar structure despite originating in different systems each with their own internal architectures. Widespread adoption of FHIR for data exchange would allow for the creation of national resources and knowledge built on the common assumptions about data structure the FHIR allows.

Many of the benefits will also extend beyond the immunization space and into other areas of public health data sharing.

Policy Considerations

Most general policy considerations that exist for immunization data exchange also exist when contemplating bulk FHIR query (e.g., a focus on interoperability standards, adherence to FHIR standards and profiles, use of standard terminology, etc.) However, using **bulk FHIR queries** for immunization data involves several distinct policy considerations compared to regular immunization messaging, mainly because of the scale, data aggregation, and potential for broader use cases. Below are key factors to consider:

Data Sharing and Consent

- **Informed Consent:** Bulk queries may include data from multiple individuals, necessitating clear policies on whether individual consent is required or if aggregate use is permitted under public health exceptions.
- **Scope of Data Access:** Bulk FHIR queries access data on more than one patient, which is broader than what is typically shared in single immunization messages. Policies must address whether access is restricted to specific use cases, like public health reporting or quality improvement.

Patient Identification

- Develop policies on how patients will be uniquely identified in bulk queries to ensure the correct matching of immunization records to the right individuals.
- Consider the implications of using demographic query by parameter approaches versus unique identifiers and the necessity of existing consent for identification.

Data Privacy and Security

- **Volume of Data:** Bulk FHIR queries retrieve large datasets, increasing the risk of exposure in case of data breaches compared to individual immunization messages.
- **Access Control:** Organizations must ensure that only authorized entities access bulk data. This might involve role-based access control or additional permissions.
- **Possible Need for De-identification:** For bulk queries, especially for research or population health, policies may require data de-identification or anonymization, unlike individual messaging, which typically uses identifiable data.

Confidentiality, Compliance and Legal Frameworks

- **HIPAA, FERPA, GDPR, Federal Rules:** Bulk queries must meet stringent compliance standards that may be specific to the entity sending the bulk query, ensuring compliance with regulations such as HIPAA (Health Insurance Portability and Accountability Act) or FERPA (Family Educational Rights and Privacy Act) in the US or GDPR (General Data Protection Regulation) in the EU, especially regarding patient consent and data sharing.
- **Public Health Statute and Rule, Public Health Reporting:** Policies may allow or disallow bulk data access specifically for public health purposes, requiring clear boundaries between permissible and impermissible uses. Differing laws may add complexity, particularly if they involve data across jurisdictions with disparate data exchange laws.
- **Data Retention:** Data retention policies may differ for regular immunization messaging and bulk queries.

Technical Infrastructure and Interoperability

- **Performance Impacts:** Bulk queries may demand significant resources from FHIR servers. Policies might need to address rate limiting or scheduling to prevent disruption of other production services. Draft guidelines for handling large data sets effectively, including limits on the size and frequency of bulk queries to avoid system overload or

performance issues. Consider policies for prioritization of queries based on urgency or public health need.

- **Data Quality:** Bulk queries might highlight discrepancies or gaps in data (e.g., missing immunization records), or might inadvertently propagate problematic data, requiring processes for reconciling such issues.

Data Usage and Secondary Analysis

- **Purpose Limitation:** Bulk FHIR data might be repurposed for analytics or research. Policies must ensure alignment with the original purpose of data collection and impose restrictions on secondary uses.
- **Equity in Data Use:** Bulk data could disproportionately benefit well-resourced organizations. Policies should address equitable access to aggregated immunization data for public health benefit.

Governance and Oversight

- **Auditing:** Policies for bulk FHIR queries may require enhanced auditing to track who accessed data, when, and for what purpose (purpose in particular may be challenging to ascertain through bulk queries).
- **Partner Engagement:** Public trust is critical; governance frameworks might need to include a broad set of community stakeholders when creating policies for bulk data use.

Integration and Harmonization of Data Sets

- **Consistent Interoperability Challenges:** Bulk FHIR queries could change the cadence of immunization exchange across partners within the health ecosystem (IIS, EHRs, HIEs, etc.), necessitating policies for harmonizing these datasets and avoiding data gaps or duplication.
- **Updates and Synchronization:** Policies should address how frequently bulk data can be queried and how updates to immunization records are synchronized.

In summary, there are no policy hard stops when considering bulk FHIR query use cases, but this alternate method of data exchange requires comprehensive policies that account for the greater data volume, broader access implications, and the potential for secondary uses beyond individual single-message immunization exchange.

Technical Concepts

Bulk Data Access FHIR IG Overview

The [FHIR Bulk Data Access Implementation Guide](#) (hereafter referred to as the Bulk Data IG) was developed and published by the [SMART Health IT project](#) at Boston Children's Hospital Computational Health Informatics Program. This standard addresses the "need to efficiently

access large volumes of information on a group of individuals”¹. As of the time of publication of this document, the latest version of the implementation guide is version 2.0.0. All efforts undertaken by the Helios Bulk Data Priority Area have focused on this version of the standard.

As described in the [Bulk Data Export Operation Request Flow section](#) of the Bulk Data IG, a bulk data exchange typically involves two primary roles, the Bulk Data Provider and the Bulk Data Client. We refer readers to the Request Flow section noted above for a detailed description of the system interactions involved. For the purposes of this document, we will refer to this sequence of interactions as the “bulk data query”. As well, for discussion of the primary immunization use case described in this document, the IIS implemented by the jurisdictional immunization program plays the role of the Bulk Data Provider (or Server) while the system employed by the Authorized User plays the role of the Bulk Data Client. In many of our examples where the Authorized User is a healthcare provider, the system in use is expected to be an electronic health record (EHR) but may be another type of system for other user types, however we expect that the technical and workflow requirements will be substantially similar for all types of Authorized Users.

Operation Types

The Bulk Data IG defines three approaches to querying for data in the form of custom operations:

- [FHIR Bulk Data System Level Export](#)
 - FHIR Operation to export data from a FHIR server whether or not it is associated with a patient.
- [FHIR Bulk Data Group Level Export](#)
 - FHIR Operation to obtain a detailed set of FHIR resources of diverse resource types pertaining to all members of the specified Group.
- [FHIR Bulk Data Patient Level Export](#)
 - FHIR Operation to obtain a detailed set of FHIR resources of diverse resource types pertaining to all patients.

For the purposes of the immunization use case, the Bulk Data Priority Area primarily explored the use of the Group Level Export operation and much of the content of this document is predicated on the use of this approach. The System Level Export operation was far too broad for the purposes of the use case. As well, fundamental to the immunization use was the concept of cohorts of individuals appropriate to share with a given Authorized User which meant that the Patient Level Export operation was similarly too broad for the use case. Note that while the Patient Level Export does offer an experimental patient parameter which allows the client to specify a specific list of patient records to include in the response, Priority Area members felt that the potentially large size of the cohort of individuals might pose a challenge to the use of

¹ <https://hl7.org/fhir/uv/bulkdata/index.html>

the Patient Level Export operation and that focusing on the Group Level Export operation was more likely to be productive.

FHIR Group Resource

The FHIR [Group](#) resource “represents a defined collection of entities that may be discussed or acted upon collectively”². While the range of potential uses of the Group resource is broad, for the purposes of the immunization use case, the Group resource is utilized as a mechanism to identify a cohort of individuals that will be the subject of a bulk data query. The Group resource can be used in two ways:

1. As an enumerated list of members of the Group
2. A description of the qualities that Group members have (without explicitly listing individual members of the Group)

In the [Group Types](#) section of this document, we further describe the possible applications for both ways of use.

Note that the list of individuals contained within the Group are references to Patient resources contained on the FHIR Server (the same place as the Group itself).

Parameters

Each of the three operations in the Bulk Data IG define a set of input parameters which can be used to specify the output of the query response. Below, we note a few input parameters to the Group Level Export operation which our exploration has suggested may be important as part of pilot and production implementations. Readers should consult the Bulk Data IG for a description of all possible parameters for each of the three export operations.

Parameter Name	Server Usage	Client Usage	Description
_since	Required	Optional	Resources will be included in the response if their state has changed after the supplied time
_type	Optional	Optional	A string of comma-delimited FHIR resource types. The response SHALL be filtered to only include resources of the specified resource types(s).
patient	Optional	Optional	A list of Patient resource references. Data will not be returned for Patients who are not included in this list.

The _since parameter is expected to be an important element of retrieving data on individuals over time. When the cohort of individuals is likely to be queried multiple times over a substantial

² <https://www.hl7.org/fhir/group.html>

period of time, the inclusion of the `_since` parameter as part of the query will allow the exchange of only data that has changed since the last query when the parameter is valued with the date/time of the last query. This may dramatically reduce the amount of data transferred, as it will be limited to only the delta since the last query, reducing burden on both the server to supply the data and the client to process it.

The `_type` parameter will also be important as a means to ensure that the Authorized User only receives data types that they are interested in and eligible to access. In the immunization use case, this may be less critical given the narrow focus of IIS (see the [Data Payload](#) section for information on the resources likely to be exchanged) but may play a larger role for other public health uses where the data source collects a broader range of data.

The use of the patient parameter is expected to be limited to a subset of use cases where the size of the cohort individuals is relatively small and an enumerated list of individuals is feasible to include in the bulk data query.

Security Expectations

Any interoperability implementation must thoughtfully consider the security and privacy implications of data sharing and the bulk exchange of public health information is no exception. This document does not explicitly discuss security expectations surrounding the use of FHIR bulk data exchange but strongly encourages implementers to ensure appropriate access controls are implemented to ensure an appropriate level of data protection. In the immunization-specific use cases, for example, providers may choose to model their FHIR access controls on the existing controls in place for their HL7 v2 messaging as influenced by local regulation and policy.

The FHIR base standard is explicitly noted as not being a security protocol, rather, it notes that FHIR-based exchange protocols must be used in association with various security protocols documented outside of the FHIR-base standard. The Bulk Data IG indicates that providers implementing bulk data should implement OAuth 2.0 access management in accordance with the [SMART Backend Services Authorization Profile](#). Please consult the [FHIR Security](#) section of the base standard for further information on security expectations.

Technology Implementation

Support for the FHIR standard can be implemented in several ways. Implementers of a bulk data solution for public health use cases should carefully consider the technical architecture they select. This document aims to provide system architecture guidance to those who will respond to bulk data queries for immunization history data, such as state immunization programs. Because program environments and needs differ considerably with respect to existing infrastructure, usage volume and other aspects, the Helios effort doesn't seek to recommend a single "best" architecture. Instead, it intends to share architectural options along with considerations for applying them to a particular environment and needs. Implementers

developing a solution should explore existing infrastructure within their organization and collaborate with other IT initiatives to implement FHIR locally.

This guidance was developed by a Helios project team tasked with defining how an immunization information system (IIS) may need to set up an architecture to support bulk FHIR queries. The considerations below were informed by the experiences of immunization organizations currently pursuing FHIR capabilities.

FHIR Implementation

Two common approaches are implementation as either a FHIR server or a FHIR façade. Thorough discussions of both of these approaches are documented online however a brief overview is provided below. A FHIR server is described here as a processing environment (with associated data stores) designed specifically to maintain and produce FHIR resources. In contrast, a FHIR façade is described as functionality that can interact with external querying systems using FHIR (accepting FHIR bulk queries and returning FHIR resources in response) while interacting with internal systems and data stores that are not FHIR-aware to compile requested information and convert it to the FHIR format.

Whether the FHIR server or FHIR façade approach is used may depend on whether the query data store is

- maintained in the IIS's primary, non-FHIR-oriented format - requiring a FHIR façade to bridge between the FHIR and internal forms

or

- in a form that natively supports FHIR representation as actual FHIR resources or data that is mapped to a FHIR model.

Some of the pros and cons of these approaches are provided below:

- **FHIR Server**
 - Pros
 - FHIR server implemented may natively support bulk queries
 - No performance impact on the production IIS system
 - Cons
 - Transformation of data must happen as data is copied to the FHIR server
 - Note that the transfer of data may be real-time or periodic (see below for further discussion)
 - May require additional development to support other functionality (eg evaluation and forecasting) using FHIR formatted data
 - Cost associated with maintaining a copy of the data in FHIR format
- **FHIR Façade**
 - Pros

- No (or minimal) transformation or processing required outside of the façade (not all data may need to be copied)
 - No performance impact on the production IIS system depending on how data is staged (see below for further discussion)
- Cons
 - FHIR Façade must be developed and supported
 - Cost associated with maintaining a replica of the data depending on how data is staged (see below for further discussion)
 - Some IIS functionality (patient matching, evaluation/forecast etc.) may need to be replicated

Staging of Data

Related to the topic of server vs façade is the concept of the source of the data used to respond to an incoming query. The immunization data that acts as the source for responses may be managed in different ways. One fundamental decision is that of using the primary IIS data store versus a replica of the data. In other words, is the IIS's transactional data store queried directly, or is the query accessing a copy of that data? If implementing a replica approach, the frequency of data replication must also be considered. That is, does the replica data get updated continuously as the transactional data store is updated (as a parallel mirror) or copied periodically?

Considerations when selecting the method for staging data may include:

- The size of the jurisdiction's data set
- The expected frequency of data queries
- The size of data expected to be returned in response to a given query
- The importance of timeliness of data
- The performance impact on other IIS functions

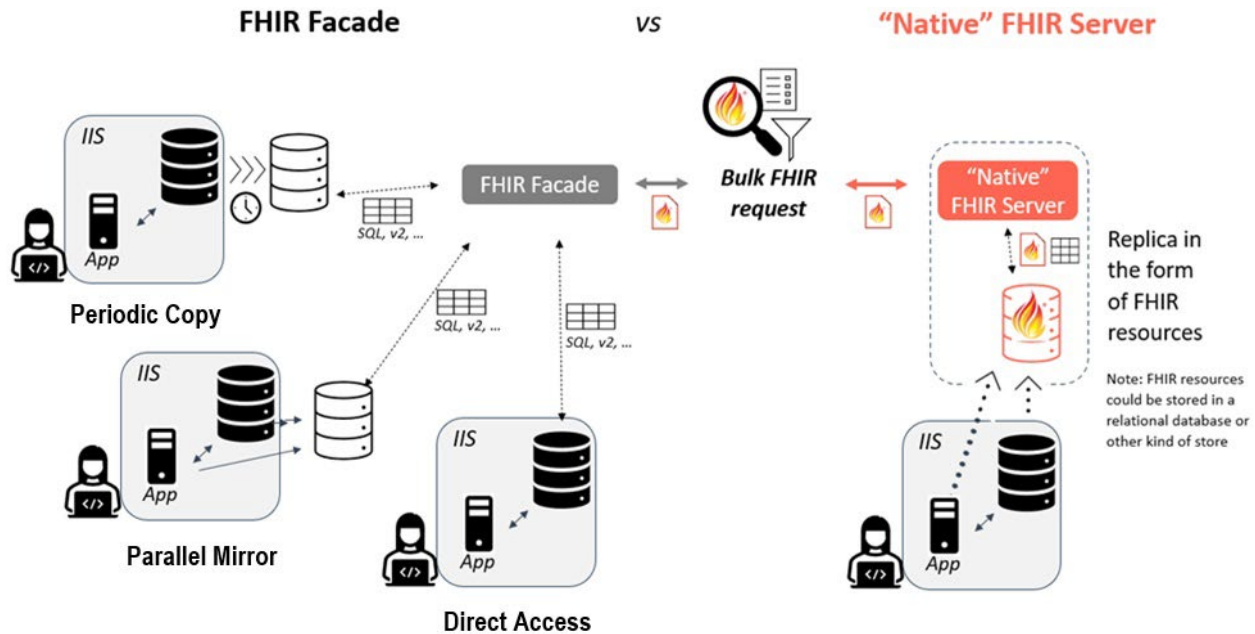
Some pros and cons to these various approaches include:

- Direct Access
 - Pros
 - Replica data doesn't need to be maintained (no additional cost)
 - Data is up to date at all times
 - No need to replicate core functionality
 - Cons
 - FHIR Façade must be developed and supported
 - Possible performance impact on the production IIS system depending on the frequency and size of queries
- Periodic Copy
 - Pros
 - No performance impact on the production IIS system
 - May be possible to use mirrored data for other analytical purposes

- Cons
 - Depending on the refresh period, data may not be fully up to date
 - Scheduling of the updates may be impacted by the other uses of the data (availability for a Data Mart may impact)
 - Cost associated with maintaining a mirror of the data
 - Some IIS functionality (patient matching, evaluation/forecast etc.) may need to be replicated
- Parallel Mirror
 - Pros
 - Data should be up to date at all times
 - No performance impact on the production IIS system
 - May be possible to use mirrored data for other analytical purposes
 - Cons
 - FHIR Façade must be developed and supported
 - Some IIS functionality (message processing, deduplication, evaluation/forecast etc.) may need to be replicated
 - Cost associated with maintaining a mirror of the data

General Considerations

The following illustrates the relationship between the concepts of FHIR implementation and staging of data.



Group Types

A Group that is the subject of a Bulk Data query may fall into one of a small number of categories depending on the mechanism used to define the group and the relative rate of

change of the members of the group. In general, a group may consist of either an enumerated list of specific individuals for a population of interest, or the group may describe the characteristics of the individuals who make up the population (without explicitly listing specific individuals to include in the Group). In the case of enumerated lists, individuals may be added to or removed from the Group as needed so that at any given point in time the list of relevant individuals is defined within the Group. In contrast, a characteristic-based Group would likely only be expanded to a list of specific individuals at the time of responding to the query. A “just in time” determination of the list of individuals may impact the query response time as the responding system must first identify the appropriate individuals which meet the characteristics of the Group before they can retrieve and make available the data for that population. A further discussion of possible group types is presented below. The optimal type of Group will depend heavily on the use case of the query and needs of the trading partners.

Persistent Groups (enumerated)

An enumerated persistent Group consists of a list of individuals where the list is expected to be relatively stable over time. That is, while new individuals may be added to the list or existing individuals removed from the list, the majority of individuals will remain associated with the Group over time. This type of group is expected to be queried multiple times over a substantial period. Example use cases include:

- A healthcare organization (HCO) creates a list of patients under their long-term care
- A payer creates a list of members they are responsible for
- A research program is tracking a defined list of individuals over a long period of time

In the [Mechanisms for Patient Identification and Group Creation](#) section below, several approaches to creating and maintaining persistent Groups are described. While many of these approaches include the possibility to remove individuals from a Group, implementers should consider how to regularly perform out of band review and maintenance of persistent Groups to ensure that the individuals included in the Group are still active and relevant. If Groups are not regularly reviewed, they may become bloated over time, impacting system performance and efficiency by inflating the data exchanges with unnecessary data.

Transient Groups (enumerated)

An enumerated transient Group consists of a list of individuals where many (or all) of the individuals are expected to change frequently or when the Group is expected to be queried for only a single time. In most use cases, the Authorized User may be the only source of information on the content of the Group, the IIS may not have any way to identify individuals to add or remove from a Group independent of the Authorized User. Example use cases may include:

- An HCO creates a list of patients scheduled to be seen the follow day, and updates the list each day with a new set of individuals

- A school district creates a list of incoming kindergarten students and executes a bulk query before the start of the school year

Note that the use cases that are expected to use relatively small enumerated transient Groups may also lend themselves to the use of the [Patient Operation](#) where the experimental patient input parameter is used to specify a list of Patient resources to include in the response as an alternative to the use of a Group.

Characteristic Groups

A characteristic Group consists of a set of characteristics which describes qualities that Group members have. Individuals in the Group may be included or excluded on the basis of these characteristics. Membership of individuals in the Group is only determined dynamically when the Group is used during the response to a query. An example use case may be when an Authorized User creates a Group which defines members as being all individuals under a specified age with immunizations recorded and who are residing in a specific zip code so that they can explore the data without the need to create an enumerated list of individuals in advance.

While characteristic based Groups may be useful for certain use cases, for the purposes of most of the remainder of this document, we will confine our discussion to the use of enumerated Groups.

Maintaining Cohorts

Once a Group has been established on the FHIR server for a particular use case serving an Authorized User, the next step is to establish an initial enumerated list of individuals belonging to the Group and subsequently maintain the membership of the group. While both Persistent and Transient Groups must be initially established, the maintenance of a Group is more relevant to Persistent Groups. As described in the [Mechanisms for Patient Identification and Group Creation](#) section of this document, there are multiple possible approaches to populating and maintaining a list of individuals in a Group. This section will describe more general considerations in maintaining a cohort of individuals.

Some considerations to keep in mind when implementing solutions to the problem of maintaining a cohort include:

- For some use cases, such as a persistent Group where all members are associated with a healthcare organization or payer, the number of individuals contained in the Group may be very large. This may impact maintenance of the group as well as response times when responding to a query for the Group.
- While Group.member.entity allows a reference to a Patient resource, the expectation is that this will refer to a Patient resource contained on the IIS FHIR server, and not a reference to a Patient resource defined on a different system (e.g., Group.member.entity should not be a reference to a Patient resource on the FHIR client).

- As described in the Technology Implementation section of this document, depending on the technical architecture of the FHIR server implementation, the list of individuals exposed in Group.member.entity may be persisted behind the scenes in a form consistent with the underlying IIS (e.g., in a relational database table or other mechanism).
- Activities impacting Patient resources in the IIS which are independent of the maintenance of the cohort of individuals defined by the Group resource may still necessitate updates to the Group.
 - As patients are merged or unmerged, the individuals in the Group may need to be added/removed/updated to reflect the records in the FHIR server and/or underlying IIS database per an agreed upon approach to patient updates.
 - For example, if two patient records are merged in the data source, if both original records are in the Group, the non-surviving record should (presumably) be removed from the Group or in the case of the creation of a net new record being created for the merged individual, the new record is added to the Group..
 - Alternatively, if a single patient record is unmerged resulting in two new patient records, if the original patient was in the Group, the original record should be removed from the Group and processes enacted to determine if either (or both) of the new records are appropriate to include in the Group.
 - As categorization of individuals happens in the IIS, for example maintaining associations between the individual and exterior organization such as a provider or school, Group memberships may need to be updated.
- Depending on the nature of the use case served by the Group, either the Authorized User system or the IIS may need to have the capacity to add/remove patients from a Group.
 - In some cases, that responsibility may fall to the IIS, for example in the case where the IIS is using its own grouping logic to create persistent Groups for a particular Authorized User organization.
 - Alternatively, the responsibility may fall to the Authorized User organization, particularly when transient Groups are being used and the Authorized User is the only system which knows who is in the Group, such as when a healthcare organization queries for individuals with upcoming appointments or a school district queries for all students in a particular grade.
 - When onboarding an Authorized User, it is crucial that both trading partners agree on which organization bears primary responsibility for maintaining the Group.

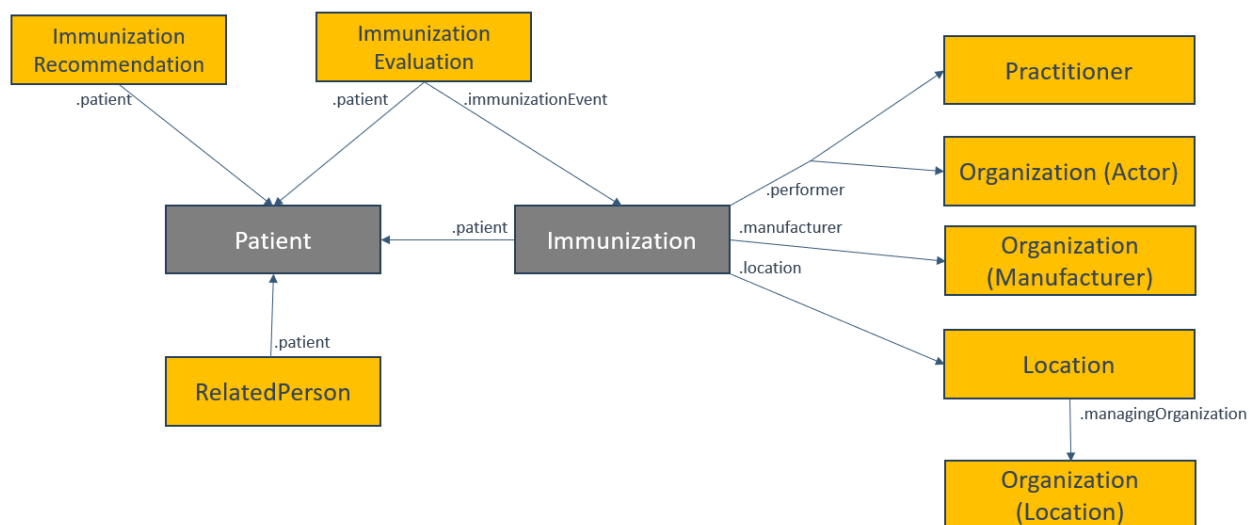
The FHIR Bulk Data Access Implementation Guide is based on [FHIR Release 4](#). Beginning in FHIR Release 5, new [Operations for Large Resources](#) are defined, this includes operations to add items to a group (the [\\$add operation](#)) and remove items from a Group (the [\\$remove operation](#)). These operations are similar to ones developed by the [Da Vinci FHIR Accelerator](#) as part of their [Member Attribution \(ATR\) List Implementation Guide](#). While neither the R5 nor ATR

operations may be directly supported by FHIR-based systems interested in Bulk FHIR, they can offer insight into implementing operations to maintain cohorts of individuals in a Group resource.

Data Payload

The data payload returned in response to a Bulk Data query can be varied. While this document does not specify the specific contents of a bulk query response to the immunization use case, below we discuss resource types which may be considered as part of the response payload. We also refer readers to the [Parameters](#) section above for an introduction to the `_type` parameter which can be used to specify the response payload.

In the case of a query to an IIS for immunization related data, a minimum of at least two resource types should be expected: Patient and Immunization. Together, these two resources should describe both the individual as well as a list of immunization events experienced by the individual as compiled by the IIS. However, responses containing FHIR Patient and Immunization resources represent the immunization history for the set of individuals, and can also include related resource types such as Practitioner, Location Organization and RelatedPerson, ImmunizationRecommendation and ImmunizationEvaluation to more fully describe the patient history and context. To align with the existing implementation and regulatory environment, adherence to US Core profiles when responding with data is highly recommended. The exact version of US Core selected may be driven by trading partner needs and other FHIR implementations of the public health jurisdiction but the selection of a version aligned with health IT regulation is advised.



Trading partners must agree upon the query response payload as part of the onboarding process to ensure that all data exchanges are appropriate for the actors involved.

Trading Partner Expectations

The [Who is Involved](#) section of this document describes 4 basic roles related to the exchange of immunization data in bulk. These are the jurisdictional immunization program, the Health Information Technology (HIT) vendor for the immunization system (AKA the IIS vendor), the Authorized User organization and the Authorized User HIT vendor. Note that references to “IIS HIT Vendor” also encompasses jurisdictions who develop and manage their own IIS locally. Each of these roles are expected to implement basic technical and/or operational functionality in support of bulk data queries regardless of the specific technology or approach selected by trading partners. These basic functionalities are described below.

Foundational Technical Requirements

The Authorized User HIT Vendor should expect to develop functionality to:

- Support one or more mechanisms to support patient matching at scale
- Create and/or update Groups
- Generate Bulk Data Access queries to immunization programs
- Retrieve and process data provided in response to a Bulk Data Access query
- Make retrieved data available to end users and/or automated processes to support organizational processes and goals

The IIS HIT Vendor should expect to develop functionality to:

- Support one or more mechanisms to support patient matching at scale
- Create and/or update Groups
- Receive and respond to Bulk Data Access queries from Authorized User organizations

Foundational Operational Requirements

The immunization program should expect to develop processes and functionality to work with the Authorized User organization to:

- Identify and document use cases for Group creation and Bulk Data query
- Create and communicate the Group IDs relevant to a given Authorized User organization
- Select appropriate approaches to patient matching, cohort creation and data exchange
 - This may include the need to identify and implement any rules or requirements that must be met in order for a patient to be added to a Group used by a particular Authorized User organization.
 - Criteria may include certainty of matching (i.e., what is a high-confidence match), establishment of a relationship between patient and organization and others
 - This evaluation may apply to a variety of methods for updating a Group including V2 mediated exchange (i.e., does a V2 message imply a relationship with an organization), flat-file mediated exchange (i.e., are all

patients in the file appropriate to add to a Group), FHIR-operations (i.e., are all Patients OK to add or remove from a Group) and more

- Identify and execute any required agreements with the Authorized User organization to exchange data on cohorts of individuals
- Implement IIS functionality required to support the agreed upon integration

The Authorized User organization should expect to develop processes and functionality to work with the immunization program to:

- Identify and document use cases for Group creation and Bulk Data query
- Select appropriate approaches to patient matching, cohort creation and data exchange
- Develop mechanisms to share Group IDs if the Authorized User organization is not permitted to directly create new Groups
- Identify and execute any required agreements with the immunization program to exchange data on cohorts of individuals
- Develop and implement processes and workflows necessary to:
 - Allow the selection of individuals for patient matching
 - Add identified individuals to a Group as appropriate
 - Remove individuals from a Group as appropriate
 - Deal with patient matching outcomes where a single high-confidence patient match is not found

Mechanisms for Patient Identification and Group Creation

The Group Level Export operation described within the FHIR Bulk Data Access Implementation Guide presupposes the existence of a Group resource on the FHIR server which contains either an enumerated list of individuals contained within the group or the criteria necessary to identify relevant individuals. The IG is silent on how individuals are identified between systems and how the Group resources are created and maintained. The Helios Bulk Data community has identified a number of different approaches to patient identification and Group creation and maintenance. In the sections that follow, we explore a variety of different approaches to accomplish these goals. During the onboarding process, trading partners should identify one or more of these approaches which are suitable for the use cases being implemented.

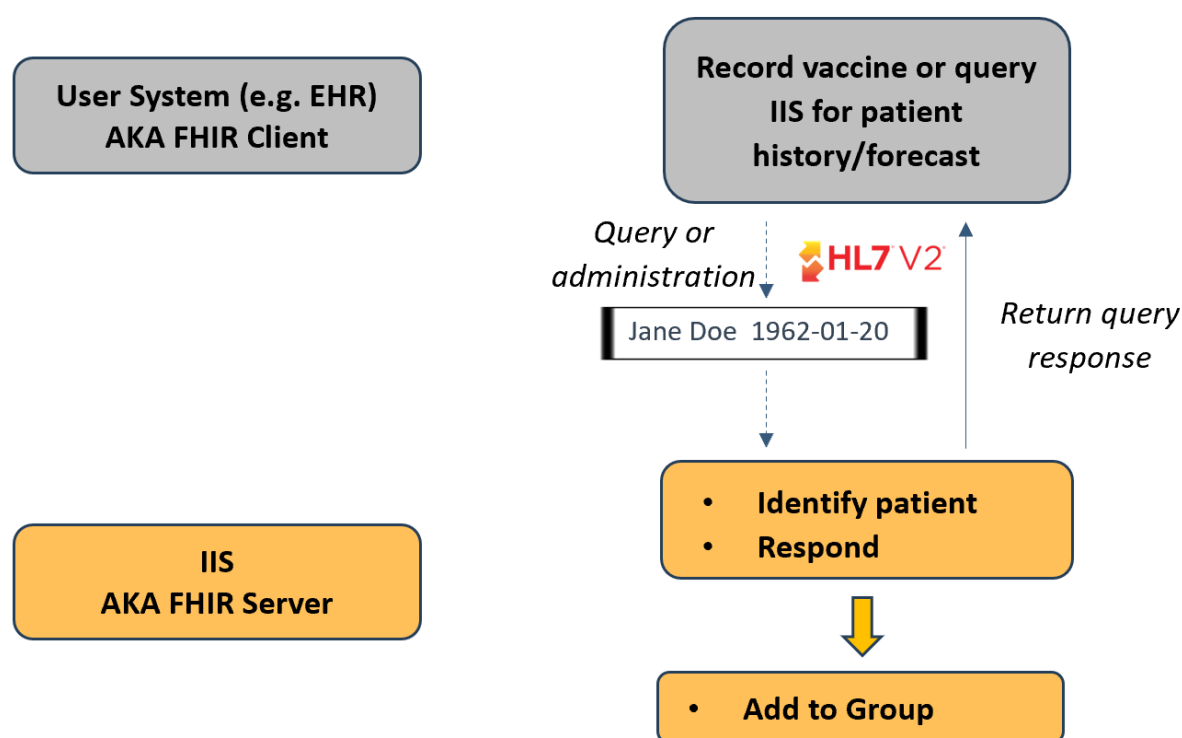
V2-Mediated Matching (Automatic Group Creation)

Description

Many trading partners, particularly healthcare organizations, have already established V2 messaging with one or more IIS. Two common exchange use cases are immunization administration event notifications (utilizing the VXU message type) and patient query for history/forecast (utilizing the QBP and RSP message types). We refer readers to the [CDC HL7 V2 Implementation Guide](#) for more details regarding the implementation of V2-based immunization exchange. Built into V2 message support by both parties are processes for

exchanging patient demographic data which can be used to match to an existing patient in the IIS. These processes can be leveraged as a prelude to FHIR-mediated Bulk Data access. The exchange of one or more V2 messages for a given patient between an Authorized User to the IIS may imply a relationship between the individual and the Authorized User. Once patient matching logic is executed upon receipt of the V2 message, the patient may be added by the IIS to a FHIR Group resource for use by the Authorized User as part of a Bulk FHIR query.

Use existing v2 messaging to build a persistent Group



Assumptions

The use of V2 messaging as a mechanism to generate Groups for Bulk Data Access will be functionality over and above any other actions taken by the IIS when processing V2 messages (such as data quality analysis, other reporting purposes, creating groups for other purposes, etc).

The FHIR Group resource to which patients are added exists for the purpose of being used as part of a Bulk FHIR query for the entire population of individuals that have a relationship with the

Authorized User. Other groupings of individuals (either using FHIR Group resources or other structures) may also exist for other purposes.

Typically, the Group ID will be shared with the Authorized User out of band (e.g., as part of the onboarding process) and will not be included in any V2 response message (ACK or RSP). Alternatively, the Group ID could be considered a non-unique patient ID in an RSP response message and stored with the patient in the Authorized User system to facilitate the use of Groups during the bulk query phase.

Group Types Supported

Persistent Groups (enumerated)

The creation of Groups via existing V2 messaging lends itself to the creation of a singular Group of enumerated individuals who are associated with the Authorized Users by virtue of either having received a vaccine from the Authorized User (thus generating a VXU message) or having been queried for by the Authorized User (thus generating a QBP message) as appropriate for local jurisdictional policies. As new individuals are the subject of a V2 message, they can be incrementally added to the Group to generate a persistent list of individuals known by the Authorized User. Note that some jurisdictions may not establish a patient care association between an individual and a provider simply on the basis of receipt of a QBP message and thus may not choose to add the individual to an Authorized User's Group based on a v2 exchange of this nature.

Benefits

The V2-mediated approach reduces the implementation lift for actors as it builds off the established V2 messaging framework. Existing patient matching logic is utilized eliminating the need to implement matching logic solely to support bulk data access. Furthermore, many IIS already establish relationships between individuals and Authorized Users in whole or in part based on v2 messaging. Many Authorized Users, primarily care providing organizations, may already be exchanging v2 messages with the IIS, reducing the complexity of the onboarding process and testing.

Limitations

A V2-mediated approach will only work with organizations who have an independent use case for V2 messaging with an IIS, and thus this may not be practical for some types of Authorized Users such as schools or other community-based organizations.

This approach does not include a direct method to remove an individual from an existing Group but rather can only add individuals to a group as a V2 message is processed. An exception may be if the individual can only have a single association per local policy and is thus removed from an existing Group before being added to a new one. Note that such an approach may impact an Authorized User's ability to receive data on all of their patients if the individual receives services from more than one provider.

For the reasons above, the V2-mediated approach is not well-suited to supporting the creation of transient Groups where the make-up of the Group changes significantly over time and there is a need to dynamically add individuals based off triggers other than those that generate V2 messages or the need to actively remove individuals from an established Group.

Technical Requirements

IIS HIT Vendor

In addition to supporting VXU and/or QBP/RSP V2 messaging, the IIS HIT vendor should implement the ability to determine the suitability of adding the individual identified as part of processing of the V2 message to a Group and if appropriate add the individual to the Group. Suitability may be based on a variety of factors including the certainty of the match, the nature of the Authorized User, and local policy regarding the establishment of a relationship between an individual and an organization.

Authorized User HIT Vendor

The Authorized User HIT vendor must support VXU and/or QBP/RSP V2 messaging with the IIS.

Operational Requirements

Immunization Program

The immunization program must support the implementation of real-time V2 messaging with the Authorized User if not already live in Production. They must also develop and implement rules for the determination of the suitability of adding an identified patient to a Group for the Authorized User.

Authorized User

The Authorized User must support the implementation of real-time V2 messaging with the immunization program if not already live in Production.

Implementation Considerations

When implementing V2 messaging for a new trading partner, the immunization program should inquire as to whether the Authorized User is also interested in performing Bulk FHIR queries. This V2-mediated approach need not be implemented for trading partners who do not intend to use Bulk FHIR queries. The immunization program should also inquire of existing trading partners if they are interested in performing Bulk FHIR queries and implement this functionality as appropriate.

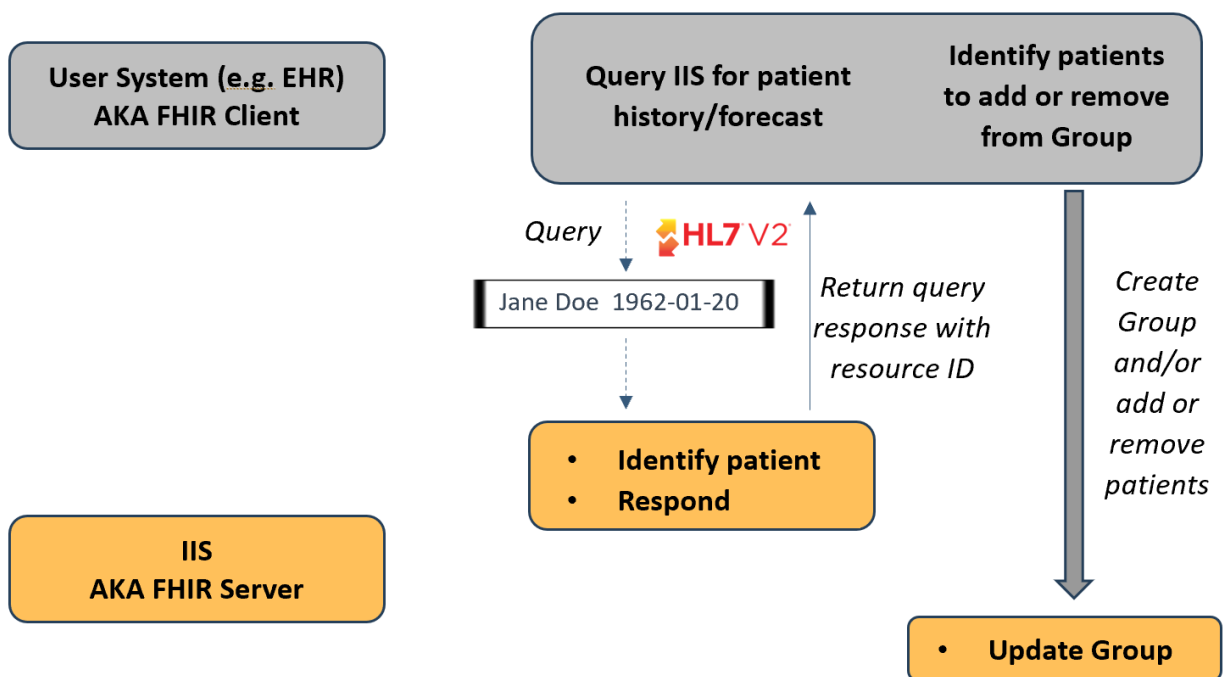
The logic as to if/when to add an individual to a specific Group based on the processing of a V2 message will be defined locally by the immunization program, including if an individual may be part of multiple Groups simultaneously.

V2-Mediated Matching (Dynamic Group Creation)

Description

Similar to the V2-mediated matching (automatic group creation) approach, existing v2 query messaging can be leveraged to support dynamic Group creation by the Authorized User. Once patient matching logic is executed upon receipt of the V2 QBP message, the FHIR Patient resource ID is returned in the RSP message as an additional patient identifier. The Authorized User system can store the Patient resource ID in association with a patient record. The Authorized User can then use the methodology described in the [Maintaining Cohorts](#) section of this document to dynamically create and/or update Groups by adding and/or removing individuals.

Use existing v2 messaging to build transient Groups



Assumptions

The use of V2 messaging as a mechanism to return the FHIR Patient resource ID for Bulk Data Access will be functionality over and above any other actions taken by the IIS when processing V2 messages (such as data quality analysis, other reporting purposes, creating groups for other purposes, etc).

Because the Authorized User system may already have an existing FHIR Patient resource ID associated with a given individual, upon receipt of a FHIR Patient resource ID in the V2 query response, the Authorized User system will validate that the returned ID is consistent with previous patient matching outcomes, for example, that the returned ID is not already associated with a different individual.

In order to ensure that the FHIR resource ID known by the Authorized User system is accurate at the time of the Bulk Data query, trading partners will agree upon how long a high-confidence patient match is valid for before a new patient match must be made.

The Group ID will be shared with the user out of band and will not be included in the V2 query response message (RSP) or the dynamic creation of Groups must be supported by both systems. Alternatively, the Group ID could be considered a non-unique patient ID in an RSP response message and stored with the patient in the Authorized User system to facilitate the use of Groups during the bulk query phase.

The decision to add or remove individuals to the Group will originate with an Authorized User system although the IIS will have the option to reject a request to add an individual if it is determined that there is no reasonable relationship between the individual and the Authorized User.

Group Types Supported

Transient Groups (enumerated)

The return of the FHIR Patient resource ID in the V2 query response allows the Authorized User system to dynamically create and/or update Groups, adding and removing patients as necessary to create transient Groups as dictated by the needs of the specific workflow being supported. For example, a transient Group of individuals with upcoming appointments can be created and then queried for.

Persistent Groups (enumerated)

The creation of Groups via existing V2 messaging lends itself to the creation of a singular Group of enumerated individuals who are associated with the Authorized Users. As new individuals are the subject of a V2 query message, the Authorized User can incrementally add the individual to an existing Group to generate a persistent list of individuals selected by the Authorized User.

Benefits

The V2-mediated approach reduces the implementation lift for actors as it builds off the established V2 messaging framework. Existing patient matching logic is utilized eliminating the need to implement matching logic solely to support bulk data access. Many Authorized Users, primarily care providing organizations may already be exchanging v2 messages with the IIS, reducing the complexity of the onboarding process and testing. Once in possession of an

individual's FHIR Patient resource ID, Authorized Users have the flexibility to make their own updates to one or more Groups allowing the creation of transient or persistent Groups.

Limitations

This approach will only work with organizations who have a use case for v2 query messaging with an IIS, and thus this approach may not be practical for some types of Authorized Users such as schools or other community-based organizations. Furthermore, because the FHIR Patient resource ID must be returned, this approach is expected to be limited to v2-based query exchanges (QBP/RSP) as the standard acknowledgment to an administration event message (VXU) does not include a standard method for returning patient identifiers.

The stability of positive patient matches is unknown. As patient demographics evolve, prior high-confidence matches may “go stale” with time. This approach does not provide a mechanism to renew the patient match short of exchanging a new v2 QBP message.

Technical Requirements

IIS HIT Vendor

In addition to supporting QBP/RSP V2 messaging, the IIS HIT vendor should implement the ability to return the FHIR resource ID for a high-confidence match in the RSP response to a QBP query message. The IIS HIT vendor must also support the ability for Authorized User system directed Group creation (optional) and the dynamic addition and removal of individuals to the Group.

Authorized User HIT Vendor

The Authorized User HIT vendor must support QBP/RSP V2 messaging with the IIS as well as the ability to store the FHIR Patient resource ID in association with the patient record upon receipt of the RSP response message. The Authorized User HIT vendor must also support the ability to support Group creation (optional) and the dynamic addition and removal of individuals to the Group.

Operational Requirements

Immunization Program

The immunization program must support the implementation of real-time QBP/RSP V2 messaging with the Authorized User if not already live in Production.

Authorized User

The Authorized User must support the implementation of real-time QBP/RSP V2 messaging with the immunization program if not already live in Production. The Authorized User must also develop and implement workflows and processes to update Groups with individuals as needed.

Implementation Considerations

Note that this approach should be implemented with a clear understanding between trading partners that V2 QBP messages should only be triggered by the Authorized User system as part of normal clinical care activities and should be used solely as a mechanism to obtain an individual's Patient FHIR resource ID. Using V2 query messages as a mechanism to obtain the resource ID could significantly strain IIS query/response capacity.

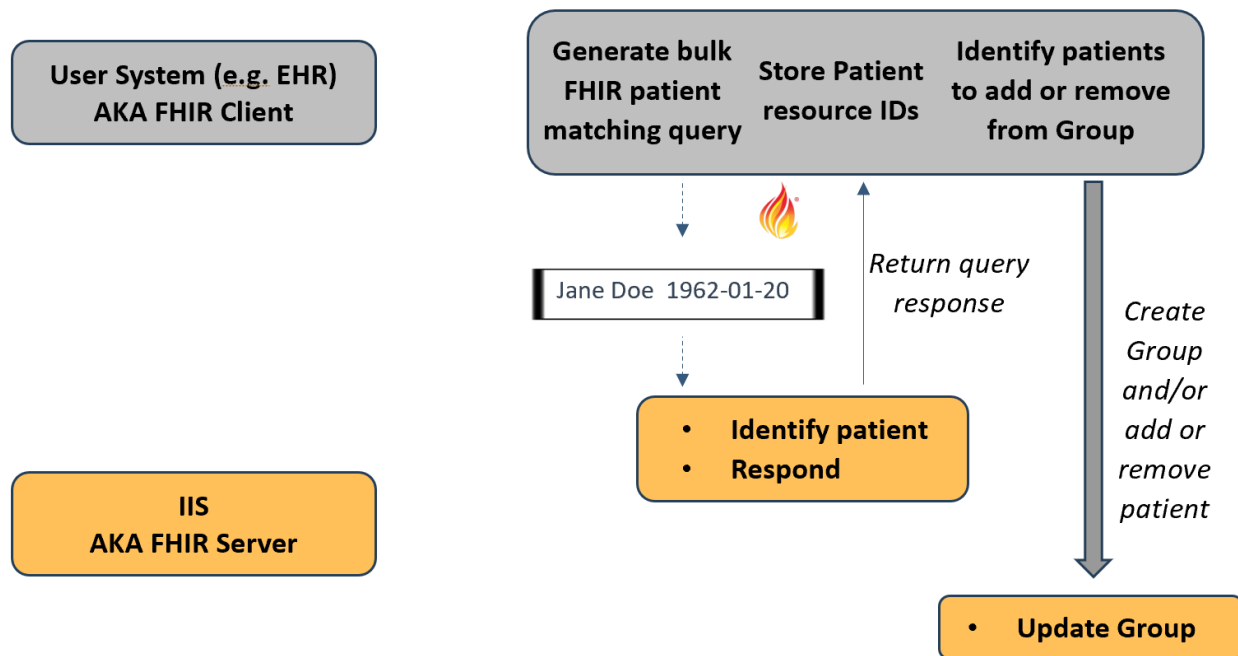
When implementing V2 messaging for a new trading partner, the immunization program should inquire as to whether the Authorized User is also interested in performing Bulk FHIR queries. This V2-mediated approach need not be implemented for trading partners who do not intend to use Bulk FHIR queries. The immunization program should also inquire of existing trading partners if they are interested in performing Bulk FHIR queries and implement this functionality as appropriate.

FHIR-Mediated Matching (Individual)

Description

The FHIR base standard supports multiple ways to identify an individual across systems. This includes both a search for a [Patient resource](#) as well as the use of the [\\$match operation](#). Both approaches include an individual's demographics as known by the Authorized User system as part of the search and can return 0, 1 or more Patient resource IDs in a response. Authorized User system can store the Patient resource ID in association with a patient record. The Authorized User can then use the methodology described in the [Maintaining Cohorts](#) section of this document to dynamically create and/or update Groups by adding and/or removing individuals.

Use Individual FHIR matching to identify patients to manage



Assumptions

The outcome of patient matching in the IIS may result in 0, 1 or multiple individuals being returned as part of the response. Authorized User systems are expected to be able to handle all possible patient matching outcomes including a single high-confidence match, multiple high-confidence matches, multiple low-confidence matches, single low-confidence match and no matches.

The Authorized User system is permitted to validate the returned Patient resource matches the input data before associating the FHIR Patient resource ID with the patient record in their system.

In order to ensure that the FHIR resource ID known by the Authorized User system is accurate at the time of the Bulk Data query, trading partners will agree upon how long a high-confidence patient match is valid for before a new patient match must be made.

The Group ID will be shared with the user out of band and will not be included in the query response. Alternatively, dynamic creation of Groups must be supported by both systems.

The decision to add or remove individuals to the Group will originate with the Authorized User system although the IIS will have the option to reject a request to add an individual if it is determined that there is no reasonable relationship between the individual and the Authorized User.

Group Types Supported

Transient Groups (enumerated)

The return of the FHIR Patient resource ID in response to a Patient search or \$match operation allows the Authorized User system to dynamically create and/or update Groups, adding and removing patients as necessary to create transient Groups as dictated by the needs of the specific workflow being supported. For example, a transient Group of individuals with upcoming appointments can be created and then queried for.

Persistent Groups (enumerated)

The identification of individuals through either a Patient search or \$match operation lends itself to the creation of a singular Group of enumerated individuals who are associated with the Authorized Users. As new individuals are identified through a FHIR Patient search or \$match operation, they can be incrementally added to the Group to generate a persistent list of individuals known by the Authorized User.

Benefits

It is expected that most FHIR enabled systems will have familiarity with a Patient search and possibly the \$match operation and offer existing functionality ready to implement.

Authorized Users have the flexibility to make their own updates to one or more Groups allowing the creation of Groups without significant burden on the immunization program.

Limitations

When the Group of individuals is large, matching individuals one-by-one can involve many interactions and may put a strain on both systems.

The stability of positive patient matches is unknown. As patient demographics evolve, prior high-confidence matches may “go stale” with time. This approach does not provide a mechanism to validate or renew the match short of performing a new search for the individual.

Technical Requirements

IIS HIT Vendor

The IIS HIT vendor must support accepting data on an individual to be matched, executing patient matching logic and returning any found matches in the FHIR format. The IIS HIT vendor

must also support the ability for Authorized User system directed Group creation (optional) and the dynamic addition and removal of individuals to the Group.

Authorized User HIT Vendor

The Authorized User HIT vendor must support the ability for the system or an end user to identify an individual that should be added to or removed from a Group in advance of a Bulk FHIR query. The Authorized User system must have the ability to initiate patient matching through either a Patient search or the \$match operation as well as validate and associate (if appropriate) a returned FHIR Patient resource ID with the individual's record. The Authorized User HIT vendor must also support the ability to support Group creation (optional) and the dynamic addition and removal of individuals to the Group.

Operational Requirements

Immunization Program

Specific operational requirements for an immunization program have not yet been identified.

Authorized User

The Authorized User must support the implementation of workflows to allow an end user or the system itself to identify individuals to be matched and initiate that process. The Authorized User must also develop and implement workflows and processes to update Groups with individuals as needed.

Implementation Considerations

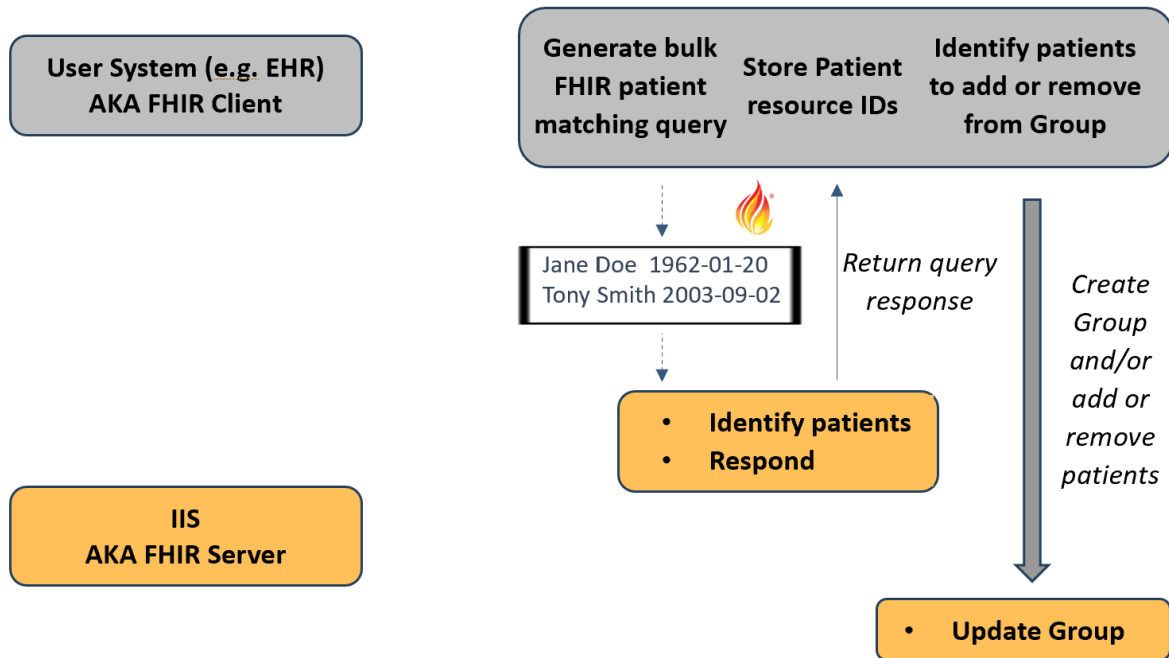
As Bulk FHIR Users are onboarded, a discussion of the system impact of doing a potentially significant number of individual queries must be considered when determining if this approach is feasible for a given implementation. The nature and size of the Groups expected to be used will be a significant factor in determining the feasibility of this approach.

FHIR-Mediated Matching (Bulk)

Description

The authors of the Bulk Data Access FHIR IG are exploring a [method to identify groups of individuals across systems](#). Their approach is to describe an asynchronous \$bulk-match operation based on the standard [\\$match operation](#). The response includes sets of matching Patient resources. The Authorized User system can store the Patient resource ID in association with a patient record. The Authorized User can then use the methodology described in the [Maintaining Cohorts](#) section of this document to dynamically create and/or update Groups by adding and/or removing individuals.

Use bulk FHIR matching to identify patients to manage



Assumptions

The outcome of patient matching in the IIS for any single individual in the bulk matching request may result in 0, 1 or multiple individuals being returned as part of the response. Authorized User systems are expected to be able to handle all possible patient matching outcomes including a single high-confidence match, multiple high-confidence matches, multiple low-confidence matches, single low-confidence match and no matches.

The Authorized User system is permitted to validate the returned Patient resource for an individual matches the input data before associating the FHIR Patient resource ID with the patient record.

In order to ensure that the FHIR resource ID known by the Authorized User system is accurate at the time of the Bulk Data query, trading partners will agree upon how long a high-confidence patient match is valid for before a new patient match must be made.

The Group ID will be shared with the user out of band and will not be included in the query response. Alternatively, dynamic creation of Groups must be supported by both systems.

The decision to add or remove individuals to the Group will originate with the Authorized User system although the IIS will have the option to reject a request to add an individual if it is

determined that there is no reasonable relationship between the individual and the Authorized User.

Group Types Supported

Transient Groups (enumerated)

The return of the FHIR Patient resource ID in response to the \$bulk-match operation allows the Authorized User system to dynamically create and/or update Groups, adding and removing patients as necessary to create transient Groups as dictated by the needs of the specific workflow being supported. For example, a transient Group of individuals with upcoming appointments can be created and then queried for.

Persistent Groups (enumerated)

The identification of individuals through the \$bulk-match operation lends itself to the creation of a singular Group of enumerated individuals who are associated with the Authorized Users. As new individuals are identified through a FHIR bulk match, they can be incrementally added to the Group to generate a persistent list of individuals known by the Authorized User.

Benefits

The \$bulk-match operation offers the ability to match large numbers of individuals asynchronously to minimize the technical load on the IIS.

Authorized Users have the flexibility to make their own updates to one or more Groups allowing the creation of Groups without significant burden on the immunization program.

Limitations

The \$bulk-match operation is still under development and is not likely to be widely supported.

The stability of positive patient matches is unknown. As patient demographics evolve, prior high-confidence matches may “go stale” with time. This approach does not provide a mechanism to validate or renew the match short of performing a new search for the individual.

Technical Requirements

IIS HIT Vendor

The IIS HIT vendor must support accepting data on a set of individuals to be matched, executing patient matching logic and returning any found matches in the FHIR format. The IIS HIT vendor must also support the ability for Authorized User system directed Group creation (optional) and the dynamic addition and removal of individuals to the Group.

Authorized User HIT Vendor

The Authorized User HIT vendor must support the ability for the system or an end user to identify a set of individuals that should be added to or removed from a Group in advance of a

Bulk FHIR query. The Authorized User system must have the ability to initiate bulk patient matching through the \$bulk-match operation as well as validate and associate (if appropriate) a returned FHIR Patient resource ID with the individual's record. The Authorized User HIT vendor must also support the ability to support Group creation (optional) and the dynamic addition and removal of individuals to the Group.

Operational Requirements

Immunization Program

Specific operational requirements for an immunization program have not yet been identified.

Authorized User

The Authorized User must support the implementation of workflows to allow an end user or the system itself to identify a set of individuals to be matched and initiate that process. The Authorized User must also develop and implement workflows and processes to update Groups with individuals as needed.

Implementation Considerations

As Bulk FHIR Users are onboarded, the frequency of bulk patient matching requests should be discussed. Depending on the size of the set of individuals involved, the \$bulk-match operation may be resource intensive, although the asynchronous nature of the operation may mitigate this concern somewhat.

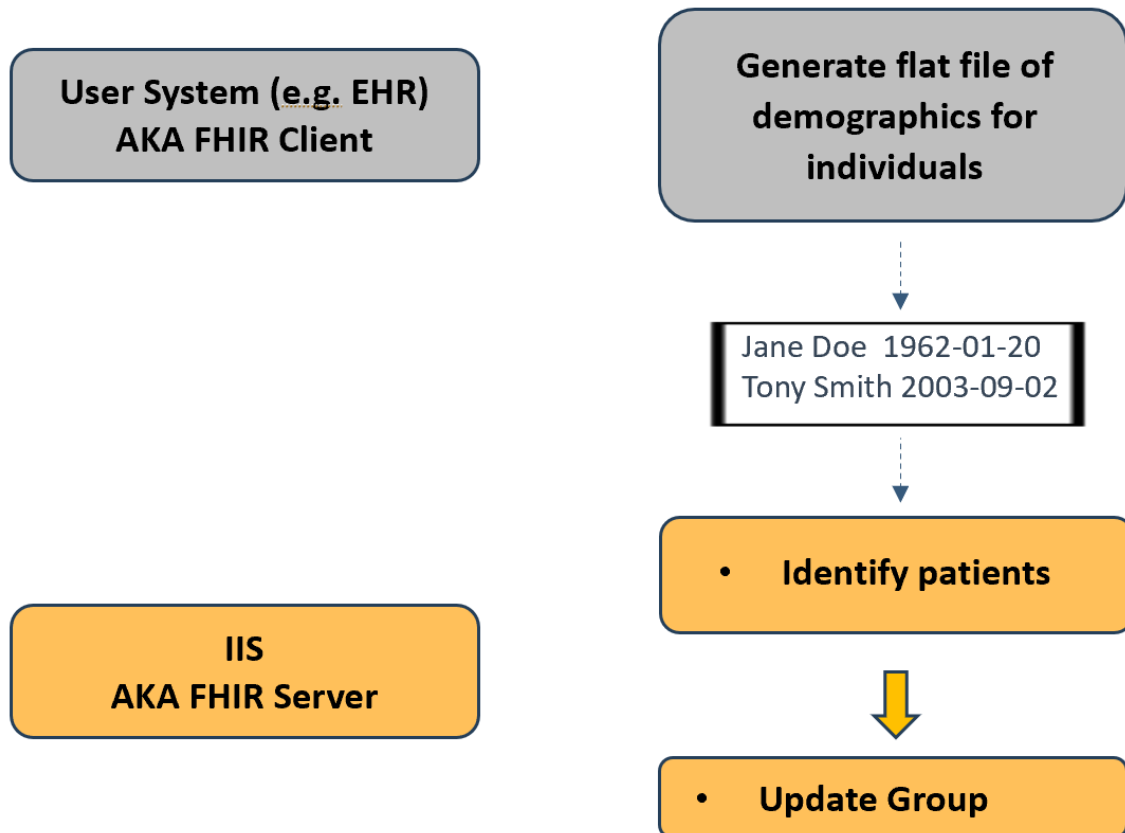
Flat File Mediated Matching

Description

Currently, bulk exchange of immunization data is often mediated by the exchange of flat files. Often, a list of individuals, including the demographic data necessary to perform patient matching, is generated by a system and shared with the immunization program as a file through SFTP or some other mechanism. The program then uses the flat file as an input into their system, identifying matching individuals and collecting immunization history data. This data is then formatted into a flat file and returned to the Authorized User of the data. There are currently no standards for generating either the input or output files, so often the file format is bespoke, creating unnecessary work for both trading partners. Because flat file mediated processes are supported by systems today, the use of a flat file to generate a Group for FHIR-based bulk data queries may be an attractive entry point for trading partners.

Purpose developed flat files produced by the Authorized User system can be used to share patient demographics for a set of multiple individuals for patient identification. Ingestion of the file and patient matching can be executed by the IIS and matching individuals can be added to a Group resource for later querying. The logic as to if/when to add a patient to a specific Group will be defined locally by the IIS in consultation with the Authorized User.

Use a flat file to populate a Group



Assumptions

The Authorized User system is capable of generating the flat file according to agreed upon specifications for the sole purpose of generating a FHIR Group for a Bulk FHIR query.

The FHIR Group resource exists for the sole purpose of being used as part of a Bulk FHIR query. Other groupings of individuals (either using FHIR Group resources or other structures) may also exist.

The Group ID will be shared with the Authorized User out of band.

Group Types Supported

Transient Groups (enumerated)

The contents of the input flat file is left to the discretion of the Authorized User, however, the use of the flat file mediated approach may lend itself best to transient groups of individuals. For example, a transient Group of individuals with upcoming appointments can be created and then queried for. The number of individuals included in the flat file and frequency of flat file submission may vary.

Persistent Groups (enumerated)

The generation of persistent Groups may be possible, however Trading Partners must agree upon the mechanisms for adding and removing individuals from the Group using flat files. Where a new flat file is agreed to be an addendum to the existing Group, as new individuals are submitted in a flat file, they can be incrementally added to the Group to generate a persistent list of individuals known by the Authorized User. However the removal of existing individuals in the Group may be more difficult with the flat file approach.

Benefits

Trading partners with existing workflows using flat files will be familiar with this process. For any systems which currently support flat file exchange, this approach minimizes the development required by the system.

Authorized Users have the flexibility to generate flat files as needed allowing the creation of transient Groups

Limitations

The exchange of a flat file is not supported by a standard FHIR workflow and must be negotiated locally.

Unless a standardized flat file format is agreed upon, IIS will have to adapt to the make up of the flat file provided by Authorized Users (or vice versa).

The use of flat files for patient matching represents only an incremental improvement on the current flat file mediated processes.

Technical Requirements

IIS HIT Vendor

The IIS HIT vendor must support the ability to receive and consume a flat file of demographics for multiple individuals and execute patient matching logic. They must also support the ability to add found appropriate matches to a Group.

Authorized User HIT Vendor

The Authorized User HIT vendor must support the ability to allow end users to generate a list of individuals to be included in a Group and to create and transmit a flat file containing the relevant demographics consistent with agreed upon requirements.

Operational Requirements

Immunization Program

The immunization program must develop an agreed upon process and format with the Authorized User to submit flat files.

Authorized User

The Authorized User must develop an agreed upon process and format with the immunization program to submit flat files. The Authorized User must also implement a workflow to allow end users to create lists of individuals to be included in a Group.

Implementation Considerations

During implementation, the trading partners must discuss the nature of the contents of the flat files and if each new flat file should replace or append the contents of the existing Group.


If Persistent groups are desired, the mechanisms for removing individuals from an existing Group must be discussed.

Reuse of Existing Groupings

Description

An immunization program may already have implemented functionality in their IIS to create associations between individuals and an organization for a purpose other than to perform a Bulk FHIR query. For example, an IIS may associate individuals with a provider or healthcare organization for the purposes of quality reporting. These associations may be replicated within a Group which can then be referenced in a Bulk FHIR query. Note that this approach is not asking IIS to change the nature of the associations they are tracking or use the existing sets of individuals for the purposes of a Bulk FHIR Query, only replicate the sets of individuals in FHIR Group form so that they are accessible by bulk query.

Reuse an existing grouping of individuals created for a different purpose



Assumptions

The requirements to create the original association should be significantly overlapping with the needs of the Bulk FHIR query.

The IIS may offer the ability to seed a Group based on existing associations but allow the User to subsequently modify the Group through some other mechanism.

As associations are updated through other mechanisms, the IIS will actively maintain the Group as well.

The Group ID will be shared with the Authorized User out of band.

Group Types Supported

Persistent Groups (enumerated)

Existing functionality to maintain cohorts of individuals in an IIS for functionality unrelated to bulk query is mostly likely to align best with persistent Groups that are maintained over time.

Benefits

Some IIS already maintain groups of individuals associated with an external partner that might be appropriate to use, thus the logic necessary to establish relationships should already exist and not require new development other than that necessary to replicate the associations in the form of a FHIR Group.

Limitations

Grouping logic may be implemented differently across IIS and thus introduce variability for HIT vendors or Authorized Users who operate across jurisdictions.

Existing grouping logic may not align well with Authorized User organization preference to perform queries. That is, the existing grouping logic may not be sufficient to provide Authorized Users access to all of the individuals they are responsible for.

Technical Requirements

IIS HIT Vendor

The IIS HIT Vendor must support the ability to leverage existing grouping functions to also add individuals to a FHIR Group.

Authorized User HIT Vendor

Specific operational requirements for an Authorized User HIT vendor have not yet been identified.

Operational Requirements

Immunization Program

An immunization program must implement functionality to establish cohorts of individuals suitable for associating with an Authorized User. The immunization program must also establish an understanding with an Authorized User to reuse existing grouping functionality for bulk query purposes.

Authorized User

The Authorized User must establish an understanding with an immunization program to reuse existing grouping functionality for bulk query purposes.

Implementation Considerations


When an Authorized User registers interest in Bulk FHIR queries, existing associations should be examined to see if they can be used as the basis for Group creation and maintenance.

Online Interface

Description

An IIS may offer an online portal including an interface to allow an Authorized User to search for an individual and add them to a Group for use in a subsequent query. Additional functionality should be implemented to allow users to manipulate existing Groups, for example removing individuals from a Group or resetting the Group to remove all members.

Use online interface to manually identify individuals



Assumptions

Authorized Users will manually enter sufficient demographics to positively identify an individual.

The Authorized User system will implement a mechanism to allow an end user to manually trigger a bulk query once a Group has been updated through the online interface.

The Group ID will be shared with the Authorized User out of band

Group Types Supported

Transient Groups (enumerated)

Where the size of a transient Group is manageable, an end user from the Authorized User may manually identify and add individuals to a Group. Given the nature of transient Groups, the need for the ability to remove an individual from the Group may be limited to those individuals entered in error, but the ability to reset the Group to an empty state or create a new transient Group is likely to be important.

Persistent Groups (enumerated)

Persistent Groups can be built up over time by an end user from the Authorized User. The ability to dynamically add and remove individuals will be important to maintain the Group over time.

Benefits

The use of an online interface allows end users to create and maintain cohorts of individuals without requiring patient matching or Group manipulation functionality to be added in the Authorized User system. It also ensures that a human is able to resolve multiple possible matches or low-confidence matches.

End users may already be familiar with IIS functionality for finding and accessing patient records easing the transition to creating Groups for bulk query.

Limitations

For large cohorts of individuals, the effort necessary to identify individuals and manipulate Groups may be prohibitively large.

Manual entry of demographics by the end user may be prone to errors.

The use of an online interface requires end users to leave the Authorized User system to identify individuals and create cohorts before returning to trigger the actual query.

Technical Requirements

IIS HIT Vendor

The IIS HIT Vendor must support the ability to allow end users to search for patients and add them to a Group through an online interface.

Authorized User HIT Vendor

The Authorized User HIT Vendor must support the ability for an end user to manually trigger a bulk query.

Operational Requirements

Immunization Program

The immunization program must implement an online portal for end user use.

Authorized User

The Authorized User must implement a workflow for end users to create Groups through an IIS online portal and then subsequently trigger the bulk query.

Implementation Considerations

This approach may have limited application except where the expected Group size is small or the Authorized User is not expected to be able to implement one of the other approaches.

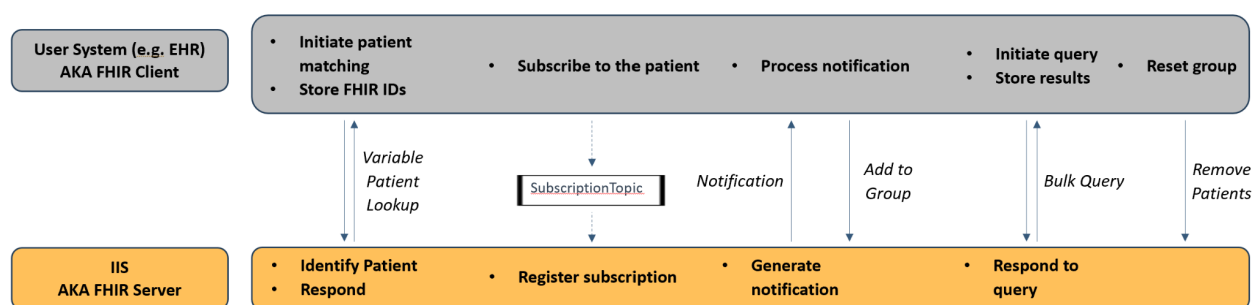
Subscription to Patient with Bulk

Description

The [FHIR Subscription](#) paradigm offers an additional mechanism for creating Groups which are limited to individuals with known updates to their immunization data in the IIS. This approach is predicated on the fact that the IIS and Authorized System have shared the FHIR Patient resource ID. This could be accomplished through any of the mechanisms described early in this document including via Bulk Patient Matching, v2 query/response, etc. Once the Authorized User is in possession of the Patient FHIR Resource ID, it subscribes to the individual via a

Subscription Topic offered by the IIS. When the individual's record is updated in the IIS, potentially including changes to either the Patient resource itself (demographics updates) or new/deleted/updated immunization events, the IIS generates a notification bundle and sends it to the Authorized User. The notification may be a robust Bundle of resources potentially including Patient, Immunization, ImmunizationRecommendation which is used to update the Authorized User system without need for further data exchange. Alternatively, the notification may only be an indication of the patient that has been updated. The Patient can then be added to a Group by the Authorized User. At a later point in time, the Authorized User initiates a Bulk Query which the IIS responds to. Upon completion of the Bulk Query, the Group resource can be reset (all current patients removed) and reused as new notifications come in.

Subscription to Patients & Bulk



Assumptions

The implementation of either a Patient SubscriptionTopic and/or Immunization.patient SubscriptionTopic will trigger a notification for patient level changes (e.g., demographics) as well as immunization event changes (doses reported, updated or deleted)

If using the Bulk Query to retrieve data, the use of the `_since` parameter may still be required (unless a full patient snapshot is desired), the value of `_since` likely corresponds to the date of the last bulk retrieval

The Group ID will be shared with the user out of band and will not be included in any communication back to the Authorized User system.

Group Types Supported

Transient Groups (enumerated)

These groups will be transient, in that they will be reset periodically, but this approach is geared towards providing updates on stable patient populations, rather than “who is being seen next week”.

Benefits

The Authorized User receives proactive notifications, possibly even in real-time, when a patient of interest is updated.

At any given point in time, the list of individuals contained in the Group are only those which have been recently updated (as determined by having been the subject of a recent notification). This should reduce the overhead associated with the query by limiting the enumerated list of Patients to only those known to have updated data.

Limitations

This approach does not solve the underlying issue of positive identification of individuals across systems.

This approach introduces new functional (FHIR subscription) that both systems must support.

Authorized Users with a large cohort of individuals will generate significant network traffic via Subscription notifications.

Technical Requirements

IIS HIT Vendor

The IIS HIT Vendor must support some mechanism of individual identification. The IIS HIT Vendor must also support the basic FHIR subscription paradigm including support for Patient and/or Immunization.subject level SubscriptionTopic and accept subscription requests from the Authorized User.

Authorized User HIT Vendor

The Authorized User HIT Vendor must support some mechanism of individual identification. The Authorized User HIT Vendor must support the ability to to subscribe to updates for a given individual and to receive and process a notification. The Authorized User HIT Vendor must also implement functionality to add individuals to a Group and to reset a Group by removing all existing Patients from the Group after a query is made.

Operational Requirements

Immunization Program

The immunization program must implement at least one mechanism for patient matching.

Authorized User

The Authorized User organization must implement at least one mechanism for patient matching as well as a workflow to trigger a Bulk Query (either scheduled or on demand).

Implementation Considerations

Trading partners will have to agree upon a mechanism for patient identification and sharing of the FHIR resource IDs, as well as the nature of the notifications sent (how robust they are).

Mechanism for Bulk Query

The mechanics of the FHIR Bulk Data Group Level Export are expected to be executed as described in the FHIR Bulk Data Access Implementation Guide. Further customization for the immunization use case is not expected to be necessary.

However, the Argonaut FHIR Accelerator, the sponsor of the IG, is exploring enhancements to the standard. The three work streams under consideration include:

- Optimization of the `_typeFilter` parameter
- Organizing the output files by Patient rather than by resource type
- Enhanced criteria based group creation

Please see the [Argonaut 2024 Bulk Optimize Project](#) site for more details.

Summary

We hope that the work of the Helios Bulk Data Priority Area as outlined in this document can serve as a launching point for immunization programs and Authorized Users to implement this functionality. We encourage all interested organizations to join the Helios Priority Area and work with us towards pilot testing and ultimately Production exchange of immunization data. Or if you have alternative use cases for other data sets, please reach out to the Leads of the Priority Area to schedule further exploration of your needs.