Journal of the American Medical Informatics Association, 30(3), 2023, 551–558
https://doi.org/10.1093/jamia/ocac237
Advance Access Publication Date: 6 December 2022
Case report

Case report

VACtrac: enhancing access immunization registry data for population outreach using the Bulk Fast Healthcare Interoperable Resource (FHIR) protocol

Leslie Lenert[1], Jeff Jacobs[2], James Agnew[3], Wei Ding[1], Katie Kirchoff[1], Duncan Weatherston[3], and Kenneth Deans[2]

[1]Biomedical Informatics Center, Medical University of South Carolina, Charleston, South Carolina, USA, [2]Health Sciences South Carolina, Columbia, South Carolina, USA and [3]Smile Digital Health, Toronto, Ontario, Canada

Corresponding Author: Leslie Lenert, MD, MS, FACP, FACMI, Biomedical Informatics Center, Medical University of South Carolina, 22 West Edge, Suite 13, Charleston, SC 29425, USA; lenert@musc.edu

Received 16 November 2022; Editorial Decision 17 November 2022; Accepted 30 November 2022

ABSTRACT

COVID-19 vaccination uptake has been suboptimal, even in high-risk populations. New approaches are needed to bring vaccination data to the groups leading outreach efforts. This article describes work to make state-level vaccination data more accessible by extending the Bulk Fast Healthcare Interoperability Resource (FHIR) standard to better support the repeated retrieval of vaccination data for coordinated outreach efforts. We also describe a corresponding low-foot-print software for population outreach that automates repeated checks of state-level immunization data and prioritizes outreach by social determinants of health. Together this software offers an integrated approach to addressing vaccination gaps. Several extensions to the Bulk FHIR protocol were needed to support bulk query of immunization records. These are described in detail. The results of a pilot study, using the outreach tool to target a population of 1500 patients are also described. The results confirmed the limitations of current patient-by-patient approach for querying state immunizations systems for population data and the feasibility of a Bulk FHIR approach.

Key words: immunization information system, population health, Bulk Fast Healthcare Interoperability Resource, case management

INTRODUCTION

COVID-19 pandemic has made the long-standing clinical problem of managing vaccine administration in target populations more apparent. Knowing who has and has not been "fully" vaccinated in a designated clinical population has always been difficult, especially when there are long gaps in care and vaccination requirements are continually changing. Further complicating the issue is the administration of vaccines in multiple sites by different clinical providers including pharmacies that are not connected electronically to population health care providers.[1] Moreover, the problem is not limited to COVID-19. Lack of access to in-person care has also created gaps in other routinely recommended vaccinations, even in previously well-managed populations.[2,3] This article describes how to newer standards might be applied to enhance both the provision and the consumption of vaccination data to further population health.

To help providers know which patients have and have not received age-appropriate vaccinations, public health jurisdictions have created Immunization Information Systems (IIS) that link and store immunization data across providers in their jurisdiction.[4] IIS are an effective intervention to improve vaccination uptake and help close vaccination gaps.[5] These systems are often driven by automatic feeds of data from electronic health records (EHRs) systems required under Meaningful Use regulations.[6]

At the individual patient level, in many advanced EHR systems, it is relatively easy to query a state IIS system to check the vaccination status of a specific patient through existing application programming interfaces (APIs) of IIS systems.[7] Population health management is more complex. Routinely checking large lists of patients, one at a time, using existing protocols, for changes in vaccination status, can create functional challenges for IISs, clinical staff, and providers. IIS interfaces to EHRs have not supported for bulk review of an entire patient panel in electronic queries. The extra work of reviewing patients one-at-a-time can negatively impact timeliness of care and efficiency of clinical staff. From a technical perspective, the number of queries for population health practice (potentially all patients, every day) is far higher than for clinical practice (patients visiting a provider for primary care or other services where vaccinations may be administered).[8] A recent monograph estimated that COVID-19-based population health activities increased the computational workloads on state IIS by as much as 10-fold, leading to performance problems.[9]

One possible solution to the problem, advocated by Lenert et al[1] is the enhancement of state IISs with the capability to perform Bulk Fast Healthcare Interoperability Resource (FHIR).[10,11] Using Bulk FHIR, providers can submit lists of thousands of patients and receive back results asynchronously, to improve the logistics of response. However, the Bulk FHIR protocol was not specifically designed for this use case and there are practical issues that may arise as a result that require refinement of the protocol. This article examines the extensions to the protocol needed to optimize Bulk FHIR query and response for immunization data, where data might come from thousands of vaccination providers and queries might come from organizations not contributing vaccination data to the IIS not contributing in an interactive fashion.

Ideally, increased capacity for reporting would be coupled increased capacity for outreach by population health providers, including smaller community organizations, to reach those residing underserved communities. A light-weight, open-source population health database tool might allow non-traditional healthcare providers, such as schools and churches, to repeatedly check the vaccination status of their students, members, etc. to maintain an accurate list for community-based efforts. This tool might also use the Bulk FHIR protocol, to synchronize the data in its system with the state immunization system or it might use a "one-at-a-time" approach to cycle through its patients if bulk query interfaces were not available. Given that there are no published tools for lightweight database systems to support bulk FHIR queries, it is an open question whether the use of Bulk FHIR for queries has sufficient advantages to warrant this new complexity. Further, this use case is different from the IIS—EHR exchange use case. In this setting, the IIS is likely to be the higher technology capability organization supporting other lower technology organizations, an unusual circumstance for public health entities.

To explore these issues, we created and tested 2 prototypes that implemented both the data request and retrieval/management process:

- An FHIR standard database that mirrors a state's Immunization Information System content and makes this content available by the Bulk FHIR standards to enhance access to IIS data. Using this system, we specifically identified potential extensions to Bulk FHIR protocols necessary to support simple and accurate bulk queries and to reduce the complexity of queries.
- A population health management program for use by providers and affiliated community organizations that can maintain a list of patients for a provider-based population-based outreach effort and communicate with a state IIS to maintain up-to-date records on the vaccination status of that population using either the Bulk FHIR protocol or repeated one-at-a-time queries.

These 2 components come together to form VACtrac, a platform designed to enhance IIS bulk response capabilities and to support both traditional and non-traditional providers' use of these capabilities in population health management.

METHODS

Figure 1 depicts the FHIR components of the VACtrac system. This system combines an FHIR immunization database that responds to bulk FHIR queries, connected to a state IIS (left most box) and a population health management program that can query the FHIR database using either bulk or one-at-a-time querying methods (middle box).

A significant focus of work was the design of a bulk FHIR server to respond to immunization queries. We achieved this by replication of the data in an IIS in the HAPI standard FHIR server (Smile CDR, Toronto, Canada). Replication of the data streams in an IIS system in an FHIR repository was achieved by both the duplication of a stream of Health Level Seven International (HL7) V2.5 vaccination messages coming into the IIS from "connected" EHRs, and by periodic export of the processed data in the IIS to the FHIR database with synchronization of data at a patient level. We believed replication of HL7 v2.X vaccination data streams might work better where a single stand-alone solution was required to enhance access to results or in conjunction with a Health Information Exchange. Replication of an IIS's database, in an FHIR server in a public health department FHIR server, might be a better option where state laws require public health control of IIS data and access. We implemented both approaches as part of testing. The replicated instance of VACtrac was designed to be used in a hybrid-cloud architecture, with the FHIR server residing in a cloud setting, to allow rapid scaling of computing resources to meet demand.

The specific use case we tested with VACtrac was monitoring of COVID-19 vaccination status of groups of patients, tracking the types of COVID-19 vaccines administered, the number of doses and the dates of vaccination, and integrating this data into the population health management portion of the program. Two sub-use cases were envisioned:

- Use case (1) access of the FHIR server by a healthcare provider that may (1a) or may not (1b) have previously submitted vaccination data to the IIS along with demographic data in the message to identify patients needing follow-up for vaccination or submitted only a partial vaccination record.
- Use case (2) access of the FHIR server by a school or university or use by an insurer to check the vaccination status of persons without a prior history of submission of vaccination record.

To implement these use cases, we used an agile development approach, testing bulk queries against the FHIR database. Based on the responses, we identified extensions necessary to improve the precision, completeness, and computational performance of responses. We worked with the authors of the HAPI FHIR server (JA and DW) and implemented extensions to the FHIR environment and deployed a new instance of the modified FHIR server in the American Immunization Registry Association sandbox environment for IIS systems, further testing the ability of the FHIR server to ingest large numbers of HL7 V2.X test messages and respond to bulk queries of individuals.

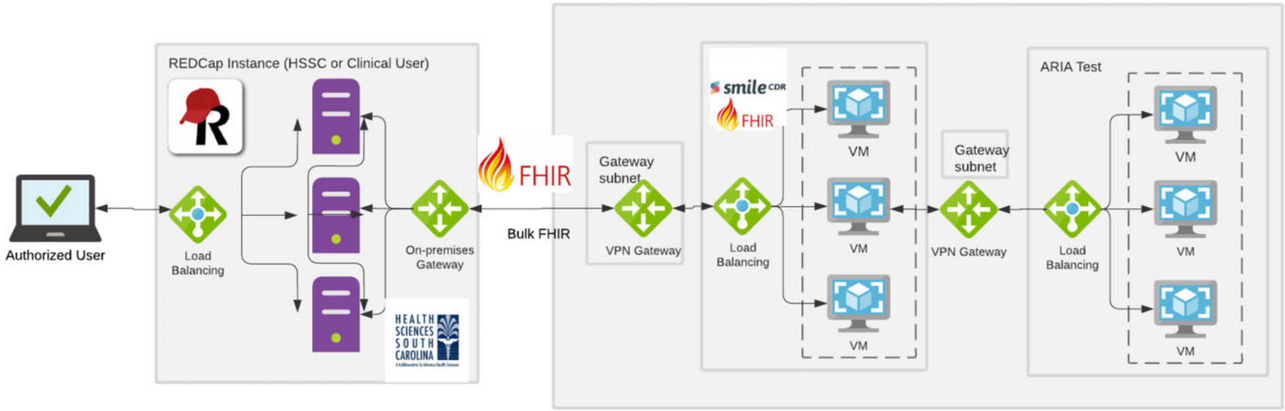

Figure 1. Architecture of VACtrac vaccination tracking tool that combined a FHIR server that replicates IIS vaccination data and a REDCap-based low-footprint population outreach tool.

Population health management platform

The population health management platform was built around a low-cost (free to academic users), highly accessible database platform that offers a secure and Health Insurance Portability and Accountability Act-compliant platform (Research Electronic Data Capture (REDCap)) and was designed to be deployed in resource-challenged settings with volunteer staff from schools, churches, or other community organizations using a web interface. This application was designed to support providers with less advanced EHRs systems that do not have population health management functionality, allowing users to upload spreadsheets with patient demographic, contact, and health data. Alternatively, VACtrac can extract clinical data for its list of patients via scheduled query of the Medical University of South Carolina's (MUSC) research data warehouse or using REDCap's Clinical Data Interoperability Services (CDIS) FHIR plug-in module for the Epic and Cerner EHRs.[12]

The REDCap database was designed to be used simultaneously by multiple providers, distributing the task of follow-up of patients across groups of case managers, with each group assigned a subset of patients for coordinated outreach. It limited access for a case manager to his or her designated set of patients at a time to preserve patients' privacy. The RECDap instance also managed clinical and non-clinical data on individual patients' risk that was used to prioritize outreach. The risk was conceptualized in 2 ways: using social determinants of health associated with patients' residential neighborhoods (census tract) and based on clinical risk factors for the progression of COVID-19 infection to severe disease. Figure 2 shows example screens from the REDCap application.

The population health tool used the standard Bulk FHIR protocol, with 3 steps: (1) posting of a query to Bulk FHIR equipped server; (2) transmission of message from the server of the availability of data from the query; and (3) downloading of the datafile prepared for the query.

To support access to IIS systems that did not have bulk FHIR capabilities, the population health tool also included software to cycle through its list of patients, one patient at a time, on a daily basis. The data flow of this application is shown in Figure 3. The application checked the IIS database using an HTTP SOAP query, with an automated login and automated cycling over its patient list. It used a variety of queues to manage waits and failures in the updating process, as shown in the figure. The resulting software was far more complex than the Bulk FHIR module, with several points of failure if the immunization registry had performance issues in responding to a query.

Population health evaluation study

We tested the components of the VACtrac system in a quality improvement study comparing different approaches for prioritizing outreach to patients. One group focused on patients residing in high social need areas. Patients' addresses were used to assign residences at the block level based on Federal Information Processing System (FIPS) codes. FIPS codes were then mapped to the JVION Vaccination Prioritization Index.[13] Patients who had undergone prior care at the Medical University of South Carolina since January 1, 2018, who had no prior record of COVID-19 vaccination in South Carolina's SIMON registry, were randomly selected for inclusion based on and assigned to 1 of 2 groups:

Group 1: residence in high, very high or highest risk areas based on the JVION index, age greater than 18 years, and no EHR-documented COVID-19 vaccine within the MUSC Health System.

Group 2: residence in very low, low, or medium risk areas based on the JVION Vaccination Prioritization Index AND one or more CDC COVID Risk factor chronic conditions[14] AND no EHR-documented COVID-19 vaccine within the MUSC Health System.

Patients were assigned to case managers by FIPS code area (so that resources could be coordinated) in blocks of 15 persons at a time, focusing on the most disadvantaged or highest illness level in each experimental group. Outreach was conducted by telephone, email, and text messaging, as supported by the REDCap application by students from the Medical University of South Carolina.

The South Carolina State Department of Health and Environmental Control chose not to participate in a no-cost replication of its SIMON database using an FHIR server. As a result, we used the patient-by-patient query mechanism for SIMON in our pilot study. This mechanism had good performance in testing in our internal American Immunization Registry Association (AIRA) sandboxes for updating of patients' vaccination status in groups of hundreds of patients. SIMON was implemented using the Envision® IIS cloud-based (Microsoft Azure®) platform and was deployed in 2018. Therefore, this study represents a "best case" environment for a one-by-one retrieval effort.

Results of the students' outreach efforts were recorded in the database, including successful contact and intent to seek vaccination.

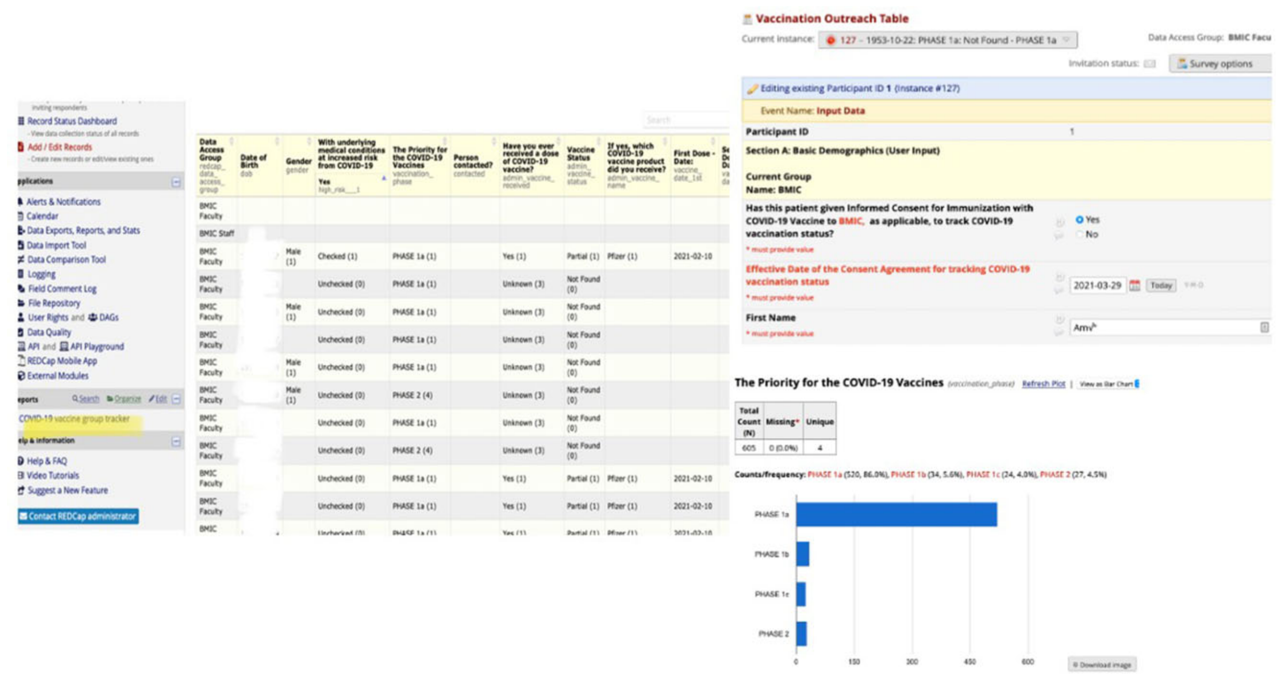

Figure 2. Selected screens from the REDCap population health management application.

Performance of the one-by-one algorithm was quantified using log files of the app to track response times and identify different potential failure types for individual and the global (1500 person) query. Examples of failure types tracked included REDCap app timing out the server query, IIS timing out REDCap session, and IIS failing to return patient data.

The overall study was classified as a quality improvement activity by the MUSC IRB.

RESULTS

Adaptations required to the Bulk FHIR Server standard

Initial work implementing Bulk FHIR Queries revealed that the specification in the Bulk FHIR standard for transmitting a complete list of patients and their identifiers for query was awkward and potentially unreliably processed at scale (thousands of persons). To address this issue, we implemented an optional Group extension to the Bulk FHIR. Instead of locally managing the group of patients and passing the set of all patient identifiers in each query, Group membership is managed on the FHIR server, and a single group identifier is passed in the query by a population health tool to the server for a query. For example, the following sample query would return the identified Pfizer COVID vaccine for all patients within group 123:

GET:
Immunization?vaccineCode=http://hl7.org/fhir/sid/cvx|300&patient._has:Group:member:_id=123

Without this group function each patient must be specified in each query, for example:

Immunization?vaccineCode=http://hl7.org/fhir/sid/cvx|300&patient.identifier=[identifier of patient 1 on the list]
Immunization?vaccineCode=http://hl7.org/fhir/sid/cvx|300&patient.identifier=[identifier of patient 2 on the list]
Immunization?vaccineCode=http://hl7.org/fhir/sid/cvx|300&patient.identifier=[identifier of patient 3 on the list]
. . . to patient n

However, this does require the FHIR server to support creation and maintenance of Groups, a complexity many IIS administrators have not heretofore managed.

In addition, we found a need to expand the parameters of the Bulk FHIR GET query. For example, a commonly used parameter in Bulk FHIR queries is "_since" parameter. This retrieves records after a time stamp to refresh the population health systems records. The extension we developed and tested was the addition of a "_mdm" parameter to the query. MDM stands for Master Data Management and is the tool in the HAPI FHIR server that replaces prior electronic master person index (e-MPI) functionality with a general approach designed to deduplicate different kinds of data types based on unique data identifiers.

The need for integrating patient matching into bulk queries was not considered in the Bulk FHIR standard. We found that a practical assumption in Bulk FHIR query model, namely, that the population health application has access to the unique patient identifier used by the IIS for integration of records, could not be relied upon in some of the use cases. For example, in use case 1b, if a provider has submitted records to the state IIS using a one-way data push approach, either has an HL7 V2.x message or in some ad hoc form such faxing a spreadsheet of administered vaccines, the provider may not have access to the unique identifier in the IIS uses for the patient (which could only be obtained in interactive communications with an IIS). They know their own identifier and could query on that identifier, but this would, without triggering e-MPI functions, result in the return of only records labeled with the institution's identifier. In use case 2, the community health organization (school, church, etc.) will not have submitted data to the IIS and may not have a unique identifier for each patient. The HAPI FHIR server has built in MDM functions that can be called to merge records.[15,16] These can operate in

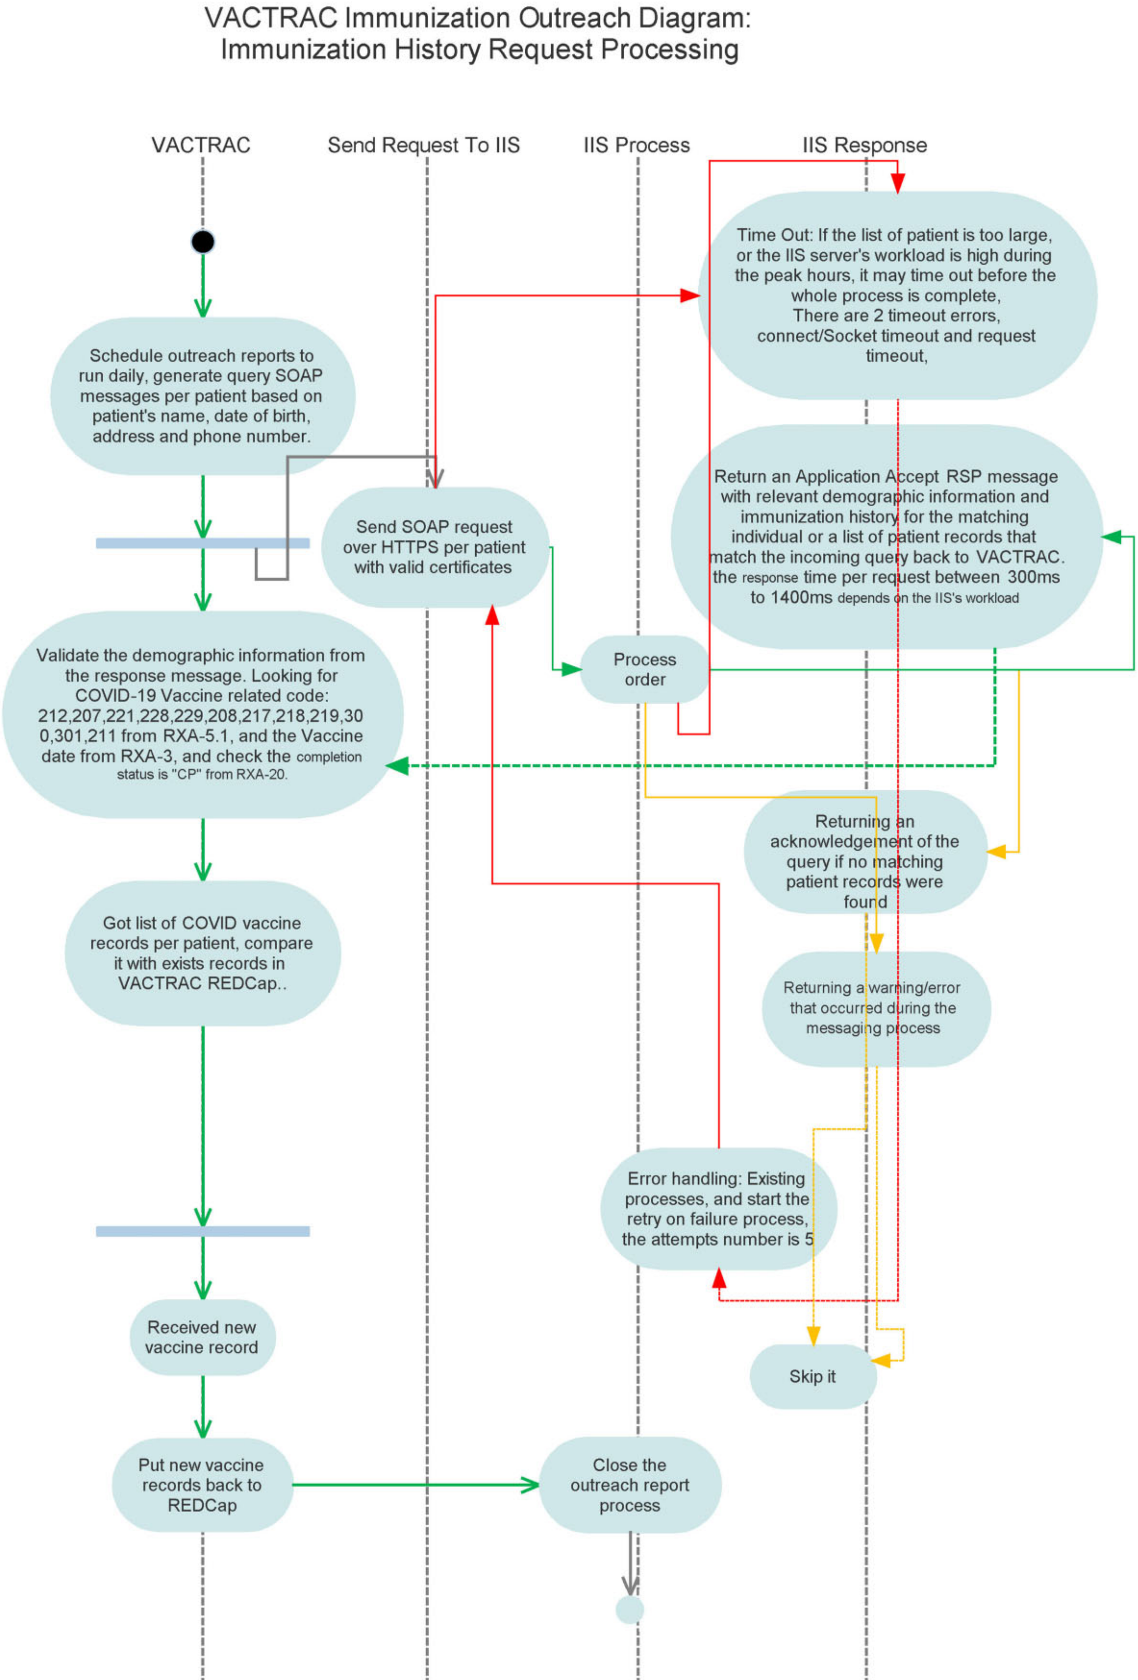

Figure 3. Data flow for a one-by-one query of an IIS from the population health application.

the background, asynchronously, replicating the algorithms used in the IIS. But, heretofore, these functions were not available in the Bulk FHIR query function.

Without this addition, at best, the population health application would need to directly query the IIS to download and store its master person index each person in the Group being monitored. This would obviate many of the advantages of a bulk approach. With the (optional) _mdm flag set matching and merging of records is automatically performed. This significantly reduces the implementation burden and preserves computing resources on the IIS side, without causing undue computational issues for the FHIR server.

However, the addition of the _mdm parameter resulted in an additional requirement to expand Group function implementation to require inclusion bundles of other FHIR resources. To use the patient matching algorithms (ie, use cases 1b and 2), the record of patients features (demographics, address, etc.) in the Group in FHIR server must be sufficiently robust to support patient matching. These data need to be uploaded by the population health application at the time of creation of a Group or the addition of a patient to the group. Fortunately, these data are likely to be available in proposed use cases and, can be readily labeled with the appropriate FHIR resource definition, allowing semantic interoperability for matching.

Scaling the Bulk FHIR system

After testing this system in an internal sandbox, with 100's of patients, we worked on scaling the system to hundreds of thousands of vaccine messages (V2.X) and tens of thousands of simulated test cases in the American Immunization Registry Association (AIRA) IIS Sandbox with connections to a cloud-based FHIR server. No further issues were identified requiring the customization of the FHIR server or other extensions to Bulk FHIR protocols.

POPULATION HEALTH TOOL

A group of 57 MUSC medical students successfully used the population health tool to conduct an outreach effort to the 1500 targeted patients. Up to 3 phone call attempts were made by the students to contact each patient totaling 2867 calls. Three hundred seventy-four patients were reached (24.9%) and 225 patients declined a conversation (15.0%). Two hundred sixty-four patients had an incorrect phone number listed in the EHR (17.6%).

The one-by-one query tool often failed to update records, with queries timing out or failing during mid-period of the pandemic. Failures occurred at all 3 postulated points: timing out of a query request on the REDCap side, timing out of a session on the IIS side, and case specification errors. Table 1 shows the errors observed at the beginning of outreach efforts (July 2021) and near the end (July 2022). Queries from the population health system decreased as patients received vaccines or opted out of further contact.

Table 1. Types of errors observed with the one-by-one method over 2 1-month periods in the pandemic

Tracking period	Mid pandemic (7/21)	Late Pandemic (7/22)
Total queries by VACtrac	48 050	30 037
Server time out errors	813	5
Session time out errors	121	1
Case specification error	1729	220
Average query response time	932 ms	391 ms
Percent queries failing	5.54%	0.75%

DISCUSSION

Access to data from immunization registries for population health activities is a topic of increasing importance and a focus of the Center for Disease Control and Prevention's new HELIOS initiative.[17] This application describes initial work to extend the Bulk FHIR protocol to address population vaccine management performed prior to the Helios initiative. This study may help inform those collaborative efforts in the HELIOS FHIR Accelerator. Two of the key issues in this FHIR Accelerator are whether IIS systems support definition of groups for bulk queries and how a provider should initiate repeated bulk queries for populations.[18]

As was learned in our agile development cycles, the most significant issue with the processing of a bulk query within an FHIR server was the need for prior access to the IIS's unique patient identifier to use an unmodified Bulk FHIR Query. Vaccination data can come from a variety of providers. Most IIS systems have some form of master-person-index solution.[4] However, a provider organization seeking access to IIS data may not have access to the IIS patient identifier, creating the need for matching algorithms *embedded* in the bulk query algorithm. The need for patient matching resulted in a requirement for sufficient demographic data to implement probabilistic algorithms, further emphasizing the need for the use of FHIR Group functionality in bulk queries.

The proposed modifications to Bulk FHIR queries have been implemented as optional functions and parameter in versions of the (open source) HAPI FHIR server published after completion of our study. As a result, developers can choose or choose not to use the query models used in this article. Our modifications reduce complexity for the users of bulk FHIR queries for immunization data, without special effort on the part of FHIR server developers for IIS systems, if implemented using open source HAPI FHIR engine, potentially speeding dissemination.

The question of the overall benefits of moving IIS systems to Bulk FHIR for population health queries is still open and is being debated as part of the Helios initiative. Table 2 contrasts 3 potential approaches: (1) augmentation of IIS systems with a Bulk FHIR engine; (2) enhancement of existing IIS infrastructure to support bulk queries: and (3) using a state Health Information Exchange (HIE) to respond to bulk queries. While there are pluses and minuses to each approach, as summarized in the table, the benefits of Bulk FHIR standards are probably greatest in the setting where a population health provider works in a border area of one or more states and has to query 2 or more different IIS systems. Here a standards-based approach make the query of multiple IIS systems, across state borders, much more feasible. Bulk FHIR may make it easier to define groups of patients for follow-up, again by offering standards, and for merging records that crosses states that are flexible and powerful. Also, by providing one uniform approach for de-duplication across entities defined using HAPI FHIR embedded algorithms may avoid conflicting results. An additional benefit of the Bulk FHIR query method is computational, allowing asynchronous query and reporting, which can preserve limited computing resources for interactive queries for use within the context of one-on-one care of a patient, and for the application of vaccine forecasting algorithms.

The availability of tools that use Bulk FHIR protocols to query patient data for population health management is also limited. Therefore, we developed a REDCap instance to manage bulk

Table 2. Comparison of options for bulk query of IIS systems

	Bulk FHIR IIS augmentation	Existing IIS enhancement for bulk query	**HIE access for bulk query**
Political barriers	Cost of replication of IIS for bulk query and computing resources	Cost of IIS enhancements and computing resources	Willingness and legal authority for public health to collaborate with the HIE. Cost of HIE enhancement if needed.
Legal authority	Existing IIS authorization	Existing IIS authorization	Legislation or policy that permits HIE to redistribute IIS data for population health practice
Governance for access to data	Manual process	Manual process	Manual process
Query group definition and maintenance	Standards-based group definition and maintenance; automated record-merging using HAPI FHIR MDM	*Ad hoc* group definition and maintenance (need to return the IIS's unique identifiers to users for cross-record queries)	*Ad hoc* group definition; automated record-merging using proprietary approaches
Query and response language	Bulk FHIR standard	*Ad hoc*	*Ad hoc*
Performance in high demand settings	Asynchronous, on-demand, optimized	Potential issues with synchronous queries on demand at high query loads	Potential issues with synchronous queries on demand at high query loads
Cross state/jurisdiction queries	Simplified by standards	*Ad hoc*	*Ad hoc*

queries. This is, to our knowledge, the first implementation of a Bulk FHIR query tool for vaccine data. However, due to political considerations, which highlighted the not-always-collaborative relationship between a state HIE and its public health authority, our test had to use a one-at-a-time query approach of the state IIS. The one-at-a-time implementation initially had a high query failure rate (5%), resulting in the data on vaccination status on the entire population not being updated. Given the adequate prior performance of the one-by-one approach in test beds and during periods of presumptive lower use, later in the pandemic, in a state-of-the-art cloud-based IIS, these results highlight the need for asynchronous bulk query capacity in IISs to respond to population health queries during periods of public health emergencies.

Limitations

This study did not test the Bulk FHIR protocol in a live clinical application. Instead, it was tested among several testbeds, while confirming prior reports of performance problems with the alternative one-at-a-time bulk query methods in a moderately sized patient population. Ultimately, the acceptability of advanced tools such as FHIR servers to IIS systems will depend on funding to states to acquire such systems and the availability of skilled personnel to use and maintain them.

CONCLUSION

With extensions, and with integration into population health management tools, Bulk FHIR may be an important tool to help providers work with IIS data to address vaccination gaps. Bulk FHIR also may make it more possible to develop agile low footprint "apps" for community organization to address vaccination gaps by reducing the complexity of interoperability for these applications.

FUNDING

This publication was supported, in part, by the National Center for Advancing Translational Sciences of the National Institutes of Health under grant number UL1 TR001450. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

AUTHOR CONTRIBUTIONS

LL co-conceived of the project, organized the effort, and wrote major portions of the manuscript. JJ, JA, WD, and KK contributed to the design and development of the software and to the publication. DW and KD contributed to the organization of the project, its funding, and to its publication.

ACKNOWLEDGMENTS

We would like to thank the American Immunization Registry Association (AIRA) and particularly Nathan Bunker and Mary Beth Kurilo for helpful discussions and access to and help with the use of the AIRA IIS test bed for testing of the VACTrac system.

CONFLICT OF INTEREST STATEMENT

Two authors (JA and DW) are members of a commercial enterprise, Smile Digital Health, which sells an enhanced, supported version of the open-source HAPI FHIR server.

SOFTWARE AVAILABILITY

The HAPI FHIR server is available at http://hapifhir.io. The REDCap master file and middleware extensions for REDcap for Bulk FHIR and for one-by-one IIS query are available at (https://github.com/vactraccovid/vactrac).

REFERENCES

1. Lenert LA, Ding W, Jacobs J. Informatics for public health and health system collaboration: applications for the control of the current COVID-19 pandemic and the next one. *J Am Med Inform Assoc* 2021; 28 (8): 1807–11.
2. DeSilva MB, Haapala J, Vazquez-Benitez G, *et al.* Association of the COVID-19 pandemic with routine childhood vaccination rates and proportion up to date with vaccinations across 8 US health systems in the vaccine safety datalink. *JAMA Pediatr* 2022; 176 (1): 68–77.
3. Nolen S. Sharp drop in childhood vaccinations threatens millions of lives. *The New York Times* 2022. https://www.nytimes.com/live/2022/07/14/world/covid-19-mandates-vaccine-cases. Accessed July 15, 2022.

4. Abbott EK, Coyle R, Dayton A, *et al.* Measurement and improvement as a model to strengthen immunization information systems and overcome data gaps. *Int J Med Inform* 2021; 148: 104412.
5. Kempe A, Hurley LP, Cardemil CV, *et al.* Use of immunization information systems in primary care. *Am J Prev Med* 2017; 52 (2): 173–82.
6. Popovich M, Altstadter B, Popovich LH. Observations illustrating the use of health informatics to link public health immunization registries and pharmacies to increase adult immunization rates and improve population health outcomes. *Online J Public Health Inform* 2016; 8 (2): e185.
7. Rajamani S, Bieringer A, Wallerius S, *et al.* Direct and electronic health record access to the clinical decision support for immunizations in the Minnesota Immunization Information System: supplementary issue: use of biomedical informatics for improving vaccine uptake and adherence. *Biomed Inform Insights* 2016; 8s2: BII.S40208.
8. Bulk Query of Immunization Information System Data. HIMSS. 2022. https://www.himss.org/resources/bulk-query-immunization-information-system-data. Accessed October 10, 2022.
9. IIP Bulk Query Toolkit. HIMSS. 2022. https://www.himss.org/resources/iip-bulk-query-toolkit. Accessed October 13, 2022.
10. Mandl KD, Gottlieb D, Mandel JC, *et al.* Push button population health: the SMART/HL7 FHIR bulk data access application programming interface. *NPJ Digit Med* 2020; 3 (1): 151.
11. Jones J, Gottlieb D, Mandel JC, *et al.* A landscape survey of planned SMART/HL7 bulk FHIR data access API implementations and tools. *J Am Med Inform Assoc* 2021; 28 (6): 1284–7.
12. Cheng AC, Duda SN, Taylor R, *et al.* REDCap on FHIR: clinical data interoperability services. *J Biomed Inform* 2021; 121: 103871.
13. Jvion launches new COVID Vaccination Prioritization Index. 2021. https://www.healthcareitnews.com/news/jvion-launches-new-covid-vaccination-prioritization-index. Accessed January 29, 2021.
14. Killerby ME, Link-Gelles R, Haight SC, *et al.*; CDC COVID-19 Response Clinical Team. Characteristics associated with hospitalization among patients with COVID-19—Metropolitan Atlanta, Georgia, March-April 2020. *MMWR Morb Mortal Wkly Rep* 2020; 69 (25): 790–4.
15. MDM Getting Started—HAPI FHIR Documentation. https://hapifhir.io/hapi-fhir/docs/server_jpa_mdm/mdm.html. Accessed October 13, 2022.
16. MDM Rules—HAPI FHIR Documentation. https://hapifhir.io/hapi-fhir/docs/server_jpa_mdm/mdm_rules.html. Accessed October 13, 2022.
17. HL. HL7® Launches Helios FHIR® Accelerator for Public Health. https://blog.hl7.org/hl7-launches-helios-fhir-accelerator-for-public-health. Accessed October 10, 2022.
18. Helios FHIR Accelerator for Public Health. https://confluence.hl7.org/display/PH/Helios+FHIR+Accelerator+for+Public+Health+Home. Accessed November 11, 2022.