
Research and Applications

Interoperability with multiple Fast Healthcare Interoperability Resources (FHIR[®]) profiles and versions

Mark A. Kramer¹ and Chris Moesel²

¹Health Innovation Center, MITRE Corporation, Bedford, Massachusetts, USA and ²Open Health Solutions Department, MITRE Corporation, Bedford, Massachusetts, USA

Corresponding Author: Mark A. Kramer, PhD, Health Innovation Center, MITRE Corporation, 202 Burlington Road, Bedford, MA 01730-1420, USA; mkramer@mitre.org

Received 22 August 2022; Revised 17 December 2022; Editorial Decision 1 January 2023; Accepted 3 January 2023

ABSTRACT

Objective: To provide a method for evaluating the interoperability of Fast Healthcare Interoperability Resources (FHIR[®]) clients and servers supporting different FHIR resources, profiles, and versions, and to determine the feasibility of FHIR servers supporting multiple FHIR Implementation Guides (IGs).

Materials and Methods: A method of analysis, the FHIR Interoperability Table (FHIT), is proposed. The FHIT involves the concept of a “catchment,” the type or category of data that a profile is intended to represent. The solution first aligns sender and/or receiver profiles according to their catchments, then determines the relationship between the admittances of those profiles, and finally interprets the relationship in terms of the feasibility of data exchange.

Results: The FHIT method is demonstrated by analyzing the FHIR-based exchange between the US Core IG and the International Patient Summary IG.

Discussion: The last few years have witnessed a significant growth in Fast Healthcare Interoperability Resources (FHIR), resulting in several major versions of FHIR, hundreds of IGs, and thousands of FHIR profiles. Previous work and available tools have not fully addressed the problem of interoperability between clients and servers that support different FHIR resources, profiles, and versions.

Conclusion: Application of the proposed methodology allows interoperability problems in FHIR networks to be identified. In some cases, new profiles that resolve those conflicts can be derived, using intersections of the original profiles. There is a need for additional tools that implement the proposed method, as well as structured methods for expressing catchments in FHIR profiles.

Key words: health information exchange, health information interoperability, FHIR, HL7

Lay Summary

A method is proposed for evaluating the interoperability of systems based on the Health Level 7 Fast Healthcare Interoperability Resources (HL7[®] FHIR[®]) standard. Exchanges based on FHIR are not automatically compatible because FHIR senders and receivers can operate under different rules and potentially deploy different versions of FHIR. The proposed method can be applied early in the design process to ensure that any proposed rules governing exchanges are mutually compatible. This is the first time that such a methodology has been proposed.

BACKGROUND AND SIGNIFICANCE

Motivation

Health Level 7 Fast Healthcare Interoperability Resources (HL7® FHIR®)¹ is a popular standard for electronic exchange of health information. FHIR defines a data exchange syntax, an application programming interface, and a set of health-related data objects called “resources.” Each type of resource represents a different category of health data.

By design, FHIR can be adapted for different healthcare use cases. Customization involves *profiling*, where FHIR built-in (or “base”) resource types are modified. The overall description of how to use FHIR for a use case is then provided by an Implementation Guide (IG). The FHIR community has created hundreds of IGs and thousands of profiles, using several different FHIR versions.² Applications following different IGs may not be able to interoperate.³

Background

Testing an individual system for conformance to the FHIR specification is supported by tools such as Touchstone,⁴ Crucible,⁵ and Inferno.⁶ Unfortunately, confirming that 2 systems conform to their respective specifications does not imply that those 2 systems can interoperate.

In FHIR, base resource definitions and profiles are represented by StructureDefinitions (SDs).⁷ An SD includes element-by-element definition of data types, value sets, cardinality, and more. Profiles derive from base resource definitions by the addition of constraints (such as narrowed cardinality or restricted data types) and specification of extension structures.

The FHIR Validator⁸ can test if a resource validates against a given SD (the terms “validates against,” “fits,” and “conforms to” are used interchangeably in this article). It includes a profile comparison function that creates a report highlighting the differences between 2 profiles. The validator also creates an *intersection profile* by combining the constraints in the source profiles, defining the characteristics of resources that conform simultaneously to both profiles.

Although the significance of intersection profiles to interoperability was noted in previous studies,^{9,10} some details were overlooked; most significantly, the matching of sender and receiver SDs according to the type of information exchanged. For example, FHIR’s BodyWeight and BodyHeight profiles¹¹ have a null intersection because they are identified by different fixed codes (Observation.code), but there is no interoperability problem because they apply to different types of data. Later, we introduce the concept of *catchments* to formalize this insight.

The relationship between 2 profiles, A and B, can be represented by a Venn diagram that represents the sets of resources conforming to each profile.¹² There are 5 possible relationships (Figure 1):

- *Equivalent*, if the same set of resources fit both profiles;
- *Subset*, if all resources that fit A also fit B (not vice versa);
- *Superset*, if all resources that fit B also fit A (not vice versa);

- *Overlapping*, if some (not all) resources fit both profiles;
- *Disjoint*, if no resource fits both profiles.

These relationships also apply when comparing different versions of the same SD. Areas of overlap represent the set of resources conforming to the intersection profile of A and B.

OBJECTIVE

The goal of this article is to provide a systematic method to analyze interoperability in an FHIR ecosystem of clients and servers that support different resource types, FHIR versions, and profiles. The resulting analysis provides the ability to detect and address interoperability problems before they arise and facilitate modifications that might avert noninteroperable systems.

METHODOLOGY

Catchments

We define a profile’s *catchment* as the type or category of data the profile is intended to represent. Profiles are designed to represent certain kinds of data in the context of a jurisdiction (“realm”) and a use case (such as claims processing). For example, the US Core (USC) MedicationRequest profile is designed to represent “an order or request for both supply of the medication and the instructions for administration of the medication to a patient, and patient-reported medications.”¹⁴ The purpose of the International Patient Summary (IPS) DeviceObserverUvIps profile is to “describe a device that plays the role of observer or performer.”¹⁵ We define an *applicable profile* to be any profile whose catchment includes the information in question, within the context of a jurisdiction and use case. For example, the USC MedicationRequest profile is applicable to patient-reported medications in the United States.

We define a profile’s *admittance* as the set of FHIR resources that conform to the profile (the circles of the Venn diagrams in Figure 1). The admittance can be larger than the catchment (Figure 2) but should never be smaller; otherwise, the profile would rule out the things it is supposed to represent. Applying the profile can be conceptualized as taking information in the profile’s catchment and mapping it into the profile’s admittance. While the admittance of a profile is objectively determinable, catchments require the interpretation of profile narratives. Later, we discuss the possibility of adding structured representations of catchments to accelerate the analysis demonstrated in this article.

The challenge of matching a profile’s admittance to its catchment derives in part from the difficulty of defining accurate and complete value sets. For example, USC Implantable Device,¹⁶ uses a value set that includes both implantable and nonimplantable devices. Removing nonimplantable devices could shrink the profile’s admittance, but if an extensible binding is used, the admittance becomes essentially infinite, unless the implementation can determine whether a novel code represents an implantable device not in the original set. Today’s




Figure 1. Possible relationships between resources conforming to StructureDefinitions A and B.




Figure 2. Conceptual relationship between catchment and admittance of a profile.

algorithms lack the semantic understanding to decide this, so the typical fallback position for extensible bindings is to accept any code. Only required bindings definitively limit the admittance of a profile.

Feasibility of exchange

Consider an exchange where the SD applied by the sender is A, and that enforced by the receiver is B. The exchange is feasible if every resource that validates against A also validates against B, that is, the relationship between A and B is *equivalent* or *subset*. It is potentially infeasible if only some possible resources conform to B—the *superset* or *overlapping* cases. Exchange is infeasible if *disjoint*. Feasibility is considered on an element-by-element basis since any infeasible element makes exchanging the resource infeasible.

In determining feasibility of exchange, the first consideration is the permitted number of occurrences. If occurrences on the sender side do not overlap those permitted by the receiver, the exchange is infeasible. If there is partial overlap, the exchange is potentially infeasible. For top-level elements, occurrences correspond to minimum and maximum cardinalities. Nested elements must account for the fact that if a parent element is optional, child elements are effectively optional.

The second consideration involves the potential values of the element. If no value on the sender side is acceptable to the receiver, the exchange is infeasible. If there is a mix of acceptable and unacceptable values, it is potentially infeasible. Data types, upper and lower bounds for quantitative values, and maximum length for strings must be considered. For elements enumerated by value sets, FHIR uses 4 binding strengths.¹³ *Preferred* and *example* do not constrain potential values, while *required* restricts the value to terms in the value set. *Extensible* implies that additional values are allowed only if the concept is not covered in receiver's value set. Typically, algorithms treat extensible bindings as if they do not constrain potential values. FHIR also supports a maximum binding to use with extensible and preferred that restricts additional codes.¹⁷

Other factors include computed constraints (invariants), the order of sliced elements, whether a slice admits additional values, the use of data absent reasons, and the representation and version restrictions on referenced resources.

Assumptions must be made to account for “Must Support” (MS), since the meaning of MS is not precisely defined by FHIR. One such assumption is that even when an element is optional, the sender *must* send that element if the data are available. Under this interpretation, MS does not affect the potential number of occurrences nor the values of an element, and therefore, has no impact in compatibility analysis.

Multiple IG support

Conflicts can also occur within a single sender or receiver that supports multiple profiles. A resource-originating system must manage

the transformation of non-FHIR data into an FHIR representation that accounts for all applicable profiles, “applicable” being defined as above in terms of profile catchments. Unlike topological catchments (watersheds), profile catchments can overlap. For example, the US Military Service History (USMSH) IG¹⁸ defines a profile for military veterans, and the Personal Health Device (PHD) IG¹⁹ defines one for patients using personal health gateways.

Assuming non-FHIR data are being mapped into FHIR for exchange purposes, multiple applicable profiles present senders with a choice of which profile(s) to apply when constructing resources. One approach is to set up separate server endpoints, one for each IG, and allow the client to choose the appropriate endpoint for its purpose. This approach works even if the applicable SDs have a null intersection but requires implementing multiple servers. Another approach (not supported in FHIR) would have the receiver specify the profile(s) the sender should apply. Profile search, using the standard search parameter `_profile`, only matches resources based on the `meta.profile` element, and neither tests conformance nor influences resource construction. Another option would be to return multiple versions of the same information, conforming to different profiles, and have the client choose the version it prefers, but this seems more confusing than setting up different endpoints. The final option is for the sender to construct resources that *simultaneously* satisfy applicable profiles, using the intersection profile. For subset and superset relationships, the most constrained profile would be implemented. For overlapping profiles, the constraints would be combined. For example, to satisfy both the US Veteran and PHD Patient profiles, resources must conform to USC patient profile as required by USMSH *and* include the specific value in `meta.profile` required by PHD.

If the intersection is null, the applicable profiles cannot be satisfied simultaneously. A sender could potentially pick one of the applicable profiles, ignoring the others, but the resulting resource would fail to conform to other applicable profiles. In some cases, a null intersection of similar profiles may not indicate an actual conflict. Consider an electronic health record (EHR) that has data on an individual's tobacco use history. Suppose it is possible to derive both the current smoking status and the number of pack-years smoked from that data, creating resources that conform to different profiles. It might appear there are 2 applicable profiles for the EHR's smoking data, but the catchments of the profiles are actually different. One profile is meant to capture smoking in the present, the other in the past. The resource originator has discretion about which profile(s) to use or populate, but this does not bear on the interoperability analysis any more than supporting other distinct, nonoverlapping profiles.

FHIR versions

Resources with different FHIR versions, or profiles based on different FHIR versions, can be analyzed in the same fashion as above, first by aligning the resources supported by the sender with the corresponding resources from the receiver and then comparing each element for changes in permissible occurrences, values, etc.

When comparing different versions of an item, the terms *forward compatible* and *backward compatible* are sometimes used.²⁰ Backward compatible means that the old item is compatible with newer versions of itself (ie, every possible resource based on the older SD will validate against the newer SD). Forward compatible means the new item is compatible with older versions of itself (ie, every possible resource based on the newer SD will validate against the older

SD). In all other cases, there are no guarantees of compatibility, but it may be possible for some resources to validate under both versions; hence exchanges are potentially infeasible (see Table 1). This article steers away from these terms because they tend to foster confusion.

Interoperability table

The objective of the FHIR Interoperability Table (FHIT) is to align the SDs applied by a sender with those enforced by a receiver. The feasibility of exchange then can be determined by evaluating the relationships between aligned SDs. The FHIT applies to a single sender/receiver combination and must be repeated to analyze multiple pairings. When populating an FHIT, only SDs that are enforced should be included. These are the SDs that define the resources the sender might send and those the receiver is willing to accept. The FHIT can be used at design time to anticipate interoperability problems or to analyze existing systems.

The table consists of the following 7 columns:

1. *Base resource* column lists all the FHIR base resource types that may be transferred from the sender to the receiver.
2. *Sender SD* column lists all SDs that the sender may apply to the resource in column 1. If more than 1 SD applies, subdivide the row.
3. *Sender Catchment* contains a textual description of the type of information the SD in column 2 applies to.
4. *Receiver SD* contains the SDs (if any) that the receiver enforces for information listed in column 3. If more than 1 SD applies, the row should be subdivided.
5. *Receiver Catchment* contains a textual description of the type of information the SD in column 4 applies to.
6. *Conclusion* holds the result: feasible if the relationship between profiles is equivalent or subset, potentially infeasible if superset or overlapping, and infeasible if disjoint. If column 4 is blank, the result is infeasible.
7. *Explanation* captures the reasons for the (potential) infeasibility.

Population of the FHIT can be aided by referencing the IG's CapabilityStatement (CS).²¹ The CS resource is required of FHIR servers and allows the server to summarize its behaviors in a standard way, providing the FHIR version, operations, supported resources, and profiles that apply to each resource. If a minimum profile (`rest.resource.profile`) is given, resources must validate against one of the supported profiles (`rest.resource.supportedProfile`) or the minimum profile.²² In practice, however, the CS is not always provided, in which case the analyst must obtain this information from system specifications, system creators, or observed behavior.

Table 1. Feasibility of exchange with different versions of the same SD

Version compatibility between SDs	Sender SD version (vs receiver)	Feasibility of exchange
Backward and forward	Either older or newer	Feasible
Forward only	Older	Potentially infeasible
	Newer	Feasible
Backward only	Older	Feasible
	Newer	Potentially infeasible
Neither backward nor forward	Either older or newer	Potentially infeasible

Systems vary in terms of how strictly they enforce profile conformance. Some systems ignore profiles and validate only against base resources. This is typical of test servers²³ and intermediaries that receive, store, or forward existing FHIR resources, such as information exchanges²⁴ and clinical endpoints that display information from multiple sources. This behavior ensures no valid resource is rejected, but data may be harder to use due to uncontrolled variability. Strict enforcement is associated with resource originators, who must pattern each new resource according to some SD, and systems that use resources in contexts that demand data conformity, such as workflow processes, analytics, and performance metrics. As mentioned above, when populating the FHIT, only SDs that are strictly enforced should be included.

RESULTS

This section analyzes the compatibility between USC version 5.0.1²⁵ and IPS version 1.0.0.²⁶ The full analysis would cover 23 resource types, 43 USC profiles, and 26 IPS profiles, but for brevity, we focus on 3 resources, Patient, Condition, and MedicationRequest.

To accelerate the analysis, a script was written to compare sender and receiver SDs focused on differences that affect interoperability (as discussed under *Feasibility of Exchange*). The FHIR Validator has similar capabilities but highlights differences not relevant to interoperability analysis and does not analyze interactions between nested elements. Invariants were left out of scope of the analysis.

Example 1: US Core (sender) to International Patient Summary (receiver)

The first example considers USC as the sender and IPS as the receiver, a scenario that could arise in international pandemic response or worldwide study of rare diseases. It is assumed that sender resources conform to a USC profile and the receiver validates against corresponding IPS profile(s). The first 2 columns of the FHIT (Table 2) are populated from USC's CS.²⁷ The third column represents the catchments of profiles in column 2, populated from profile descriptions. The fourth column lists the IPS profiles that correspond to those catchments, and the fifth are the catchments of the IPS profiles.

For Patient, there is one USC profile whose catchment is all US patients. The IPS profile applies to all patients. IPS patient requires a birth date, but USC does not, so this exchange is potentially infeasible. Additionally, `contact.relationship` has different value sets, and the IPS binding is required. Where there are semantically equivalent codes (eg, N [next-of-kin] and FAMMEB [family member]), a potential workaround is to send both USC and IPS codes, but this requires the server awareness of the receiver requirements. The USC binding of `communication.language` is extensible, and the IPS binding is required, so this element is also potentially infeasible even though the USC value set is a subset of the IPS value set.

For Condition, there are 2 profiles in USC, one for encounter diagnoses and the other for problems and health concerns. There is a single IPS profile that applies to all conditions. Both exchanges are potentially infeasible for several reasons, for example, `clinicalStatus` is optional in USC but required in IPS; IPS attempts to restrict the data types of `onset[x]` and `abatement[x]`; `category` has extensible bindings to value sets with similar codes, but from different code systems.

For MedicationRequest, there is one USC profile and one IPS profile. USC's catchment includes patient-reported medications,²⁸

Table 2. Sample rows from the FHIT for US Core (sender) to IPS (receiver)

Sender base resource	Sender SD	Sender catchment	Receiver SD	Receiver catchment	Conclusion	Short explanation
Patient	US Core Patient	Patients in US realm	Patient (IPS)	Patients in universal realm	Potentially infeasible	birthDate is required in IPS but not in USC
Condition	US Core Condition Encounter Diagnosis	Encounter diagnoses	Condition (IPS)	All conditions	Potentially infeasible	clinicalStatus is required in IPS but not in USC
	US Core Condition Problems and Health Concerns	Conditions “categorized as a problem or health concern...”	Condition (IPS)	All conditions	Potentially infeasible	clinicalStatus is required in IPS but not in USC
Medication request	US Core Medication Request	All patient medications, past, present, planned, prescribed or not, including reported medications	Medication Statement (IPS)	Record of medication being taken by a patient	Infeasible	Different base resource
			Medication Request (base resource)	Prescriptions	Potentially infeasible	Feasible only if receiver supports the MedicationRequest resource

but in IPS, reported medications are represented by MedicationStatement resources. If IPS receives a MedicationRequest where reported is true (indicating indirect knowledge), it can be rejected for using the wrong base resource. In fact, all MedicationRequests could be rejected, since IPS does not require support of this resource.

Example 2: Simultaneous support of US Core and International Patient Summary

A related question is whether an FHIR sender could simultaneously implement both USC and IPS. This involves using the catchments already aligned in the FHIT and evaluating the feasibility of the intersection profiles.

As discussed, IPS patient requires a birth date, so the intersection profile requires it. For contact.relationship, the required IPS value set must be used but all terms therein qualify as valid extensions to the USC set. Although there is no conflict, terms appearing in USC but not in IPS, for example, employer, cannot be expressed. With these reservations, a common patient profile between USC and IPS is feasible.

For Condition, there are 2 intersections, corresponding to the 2 profiles in USC. Both require clinicalStatus and subject.reference, restrict the types for onset[x] and abatement[x] to dateTime or Period, and must take a category code from USC.

As discussed above, simultaneous support of patient medications is infeasible because USC and IPS use different base resources.

DISCUSSION

After 10 years of development and 4 (going on 5) major versions of FHIR, users are still being cautioned that there are “many uninvestigated issues associated with this use [of supported] profiles.”²⁹ One missing piece is the ability to define which resources should conform to which profile. This article has shown that FHIR compatibility between senders and receivers implementing different IGs cannot be based solely on the information contained in profiles. Until computable catchments are developed, human interpretation of profile narratives is necessary to determine what profiles are intended to represent.

Defining computable catchments for profiles may appear to be a circular problem, since a catchment is a region of information space, and in FHIR, regions of information space are represented as profiles. Nonetheless, there could be an algorithmic relationship between particular values and a given profile. An example is an association between a LOINC³⁰ code and an Observation profile. Ideally, the catchment would be a function that accepts information in FHIR-compatible form and returns true only if the profile applies. Potentially, FHIRPath³¹ could serve this purpose. The potential advantage of FHIRPath is that it contains functions that invoke terminology services to translate codes between code systems and to test subsumption relationships. Another potential solution is captured in an experimental FHIR R5 profile mapping extension.³² This extension defines a relationship between a search string and a profile. A resource that would fit the search string would be expected to conform to the named profile. A third possibility is to use a distinct profile to define the catchment. The catchment profile would not define the target representation of the data, but only the characteristics that would trigger the use of the associated data profile. It remains to be seen which of these proposals best mitigates the problem of profile applicability.

CONCLUSION

This article introduced a method and criteria for interoperability in FHIR systems that involve different FHIR versions, resources, and profiles. Given the growth in the number of IGs and the existence of several major FHIR versions, the need for compatibility analysis has become critical. Although tooling currently lags, once profiles are manually aligned according to their catchments, the remainder of the analysis can be automated.

To be clear, the issue of infeasibility in FHIR systems does not arise unless SDs are enforced. It is easier to exchange data if SDs are not enforced, but then exchange participants should be prepared to process resources of all kinds. The motivation behind IGs is to impose greater uniformity. If compliance is not expected or required, IGs do not add much value.

Although the significance of intersection profiles was noted in previous studies, some details were overlooked. Considering

intersection profiles without consideration of catchments does not yield useful results. For now, human input is needed to determine catchments based on narrative descriptions. In the future, FHIRPath expressions, catchment profiles, or query strings could be employed to express computable catchments in a form amenable to computation.

FUNDING

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

AUTHOR CONTRIBUTIONS

MAK: primary author and corresponding author. CM: coauthor. Both the primary author and coauthor claim substantial contributions to the conception or design of the work, drafting the work or revising it critically for important intellectual content, final approval of the version to be published, and agree to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved.

CONFLICT OF INTEREST STATEMENT

None declared.

DATA AVAILABILITY

All data are incorporated into the article.

REFERENCES

- Health Level Seven International. HL7 FHIR Release 4B. 2022. <https://hl7.org/fhir/R4B>. Accessed July 8, 2022.
- Firely. Simplifier.net: Where FHIR Profiles are Published. 2019. <https://simplifier.net/>. Accessed July 8, 2022.
- Health Level Seven International. FHIR Conformance Rules. 2022. <https://hl7.org/fhir/R4B/conformance-rules.html#2.1.0>. Accessed July 8, 2022.
- AEGIS.net, Inc. The Touchstone Platform for HL7[®] FHIR[®]. 2022. <https://www.aegis.net/touchstone/>. Accessed July 2, 2022.
- MITRE Corporation. FHIR Crucible. 2016. <https://github.com/fhir-crucible>. Accessed July 14, 2022.
- The Office of the National Coordinator for Health Information Technology. Inferno. 2022. <https://inferno.healthit.gov>. Accessed July 8, 2022.
- Health Level Seven International. FHIR Resource StructureDefinition. 2022. <https://www.hl7.org/fhir/R4B/structuredefinition.html>. Accessed July 7, 2022.
- Health Level Seven International. Using the FHIR Validator. 2022. <https://confluence.hl7.org/display/FHIR/Using+the+FHIR+Validator/>. Accessed July 11, 2022.
- Kramer MA. Interoperability of FHIR Profiles. 2016. <https://lightmyfhir.org/2016/04/05/interoperability-of-fhir-profiles/>. Accessed July 2, 2022.
- Kramer MA, Hadley M, Mathews J, *et al*. FHIR profiling strategies for interoperability. In: *16th International HL7 Interoperability Conference (IHIC 2016)*; June 13–15, 2016; Genoa, Italy.
- Health Level Seven International. Device Metric Observation Profile. 2022. <https://www.hl7.org/fhir/observation-core.html>. Accessed July 12, 2022.
- Health Level Seven International. Supporting Multiple Profiles. 2022. <https://www.hl7.org/fhir/R4B/profiling.html#mixing>. Accessed July 2, 2022.
- Health Level Seven International. FHIR Binding Strengths. 2022. <https://www.hl7.org/fhir/R4B/valueset-binding-strength.html>. Accessed July 11, 2022.
- Health Level Seven International. US Core Resource Profile: Medication List. 2022. <https://www.hl7.org/fhir/us/core/STU5.0.1/medication-list.html>. Accessed July 6, 2022.
- Health Level Seven International. International Patient Summary Profile: Device (Performer, Observer). 2022. <https://hl7.org/fhir/uv/ips/STU1/StructureDefinition-Device-observer-uv-ips.html>. Accessed July 18, 2022.
- Health Level Seven International. US Core Resource Profile: Implantable Device. 2022. <https://hl7.org/fhir/us/core/STU5.0.1/StructureDefinition-us-core-implantable-device.html>. Accessed July 20, 2022.
- Health Level Seven International. Extension: maxValSet. 2022. <http://hl7.org/fhir/R4B/extension-elementdefinition-maxvalueset.html>. Accessed November 10, 2022.
- Health Level Seven International. FHIR Implementation Guide: Military Service History and Status STU 1 Ballot. 2021. <https://hl7.org/fhir/us/military-service/2021Sep/>. Accessed July 19, 2022.
- Health Level Seven International. Personal Health Device Implementation Guide STU 1. 2022. <https://hl7.org/fhir/uv/phd/STU1/>. Accessed July 20, 2022.
- Heidel S. Backward vs. Forward Compatibility. 2020. <https://stevenheidel.medium.com/backward-vs-forward-compatibility-9c03c3db15c9>. Accessed July 8, 2022.
- Health Level Seven International. FHIR Resource CapabilityStatement. 2022. <https://www.hl7.org/fhir/R4B/capabilitystatement.html>. Accessed July 19, 2022.
- Health Level Seven International. FHIR Resource CapabilityStatement: Supported Profile Definition. 2022. <https://www.hl7.org/fhir/R4B/capabilitystatement-definitions.html#CapabilityStatement.rest.resource.supportedProfile>. Accessed July 19, 2022.
- Health Level Seven International. FHIR Public Test Servers. 2022. <https://confluence.hl7.org/display/FHIR/Public+Test+Servers>. Accessed July 13, 2022.
- Office of the National Coordinator for Health Information Technology (ONC). HealthIT.gov. 2020. <https://www.healthit.gov/topic/health-it-and-health-information-exchange-basics/what-hie/>. Accessed July 8, 2022.
- Health Level Seven International. US Core Implementation Guide Version 5.0.1. 2022. <https://www.hl7.org/fhir/us/core/STU5.0.1/>. Accessed July 12, 2022.
- Health Level Seven International. International Patient Summary Implementation Guide Version 1.0.0. 2022. <https://hl7.org/fhir/uv/ips/STU1/>. Accessed July 14, 2022.
- Health Level Seven International. US Core Server CapabilityStatement. 2022. <https://www.hl7.org/fhir/us/core/STU5.0.1/CapabilityStatement-us-core-server.html>. Accessed July 19, 2022.
- Health Level Seven International. US Core Implementation Guide Medication List Guidance. 2022. <https://www.hl7.org/fhir/us/core/STU5.0.1/medication-list.html>. Accessed July 11, 2022.
- Health Level Seven International. Profiling FHIR. 2022. <http://hl7.org/fhir/R4B/profiling.html>. Accessed November 10, 2022.
- Forrey AW, McDonald CJ, DeMoor G, *et al*. Logical observation identifier names and codes (LOINC) database: a public use set of codes and names for electronic reporting of clinical laboratory test results. *Clin Chem* 1996; 42 (1): 81–90.
- Health Level Seven International. FHIRPath Normative Release (v2.0.0). 2020. <https://hl7.org/fhirpath/N1/>. Accessed July 14, 2022.
- Health Level Seven International. FHIR Tooling Extensions IG (v0.1.0). 2022. <http://build.fhir.org/ig/FHIR/fhir-tools-ig/StructureDefinition-profile-mapping.html>. Accessed November 8, 2022.