FHIR for Clinicians

Orion Health White Paper
Dr. David Hay
Product Strategist
022017

ORION HEALTH

How to blaze through FHIR healthcare IT projects

The FHIR Standard is maturing, and increasingly being adopted for integration projects across the globe. It is becoming essential that clinicians and informed patients become involved in these projects as they depend on the correct representation of clinical information being shared between clinical systems, which is an area that is notoriously difficult to get right. In this paper we will discuss what a clinician needs to know about FHIR to be able to knowledgeably participate in these projects. We will start with discussing the fundamentals of FHIR, focusing on those aspects that the clinical user should be familiar with, and then discuss tooling that is being developed to help people gain an understanding of FHIR. Finishing with the overall process that can be used to develop the FHIR artifacts.

Fundamentals of FHIR

FHIR® or Fast Healthcare Interoperability Resources is a new IT healthcare standard. It is one of the next generation HL7® standards in healthcare data integration, and is focused on decreasing interoperability costs and unlocking technical innovation in healthcare. FHIR will help to break down information silos that exist in healthcare. For example, a lack of readily available patient health history often forces doctors in emergency departments to make educated assessments about appropriate medication, when there is no one to speak for the patient. FHIR could support access to such vital data at the push of a mobile app button. FHIR represents a major standards upgrade that will boost access to health information and support ambitious plans for an app store for the health sector. FHIR aims to speed application development and interoperability, plus boost information sharing in healthcare, especially on mobile platforms.

FHIR was started following a HL7 Board meeting in 2011, where the board noted that interoperability requirements were changing. They noted:

- There was the need for real-time access for APIs, especially with the uptake of mobile device use.
- They saw a vast increase in the amount, type, and source of data particularly with the increasing use of personal devices.
- There was a movement to include the patient in their own healthcare.
- Data from genomics and precision medicine had increased the amount of data available to be included in a patient's electronic health record.

With aging populations and the increase in chronic diseases, the ability to data mine for analytics and to be proactive with population health management was becoming critical, and better information was required. To achieve this, implementers expect a modern standard capable of managing this huge diversity and size of data. A fresh look was needed and so FHIR was born.

FHIR is built around units called "resources", the basic unit of interoperability. The smallest amount of information that can be exchanged at one time, a resource represents a 'thing' of interest in healthcare. It could be clinical such as a Problem, a Medication or an Allergy, or administrative such as an Encounter, an Appointment, a Practitioner or a Claim.

By itself, a single resource has only limited value. To exchange more meaningful information, you join resources together, as you would make a model from Lego blocks – a process called 'resource references'. The images on the next page show a simple consultation in textual format, followed by a 'graph' of resources showing how they inter-connect to represent the same information in a highly structured way. The information represented in this way, is able to be understood by a recipient, but can also be made available for other purposes such as Decision Support or Population Based Analytics.

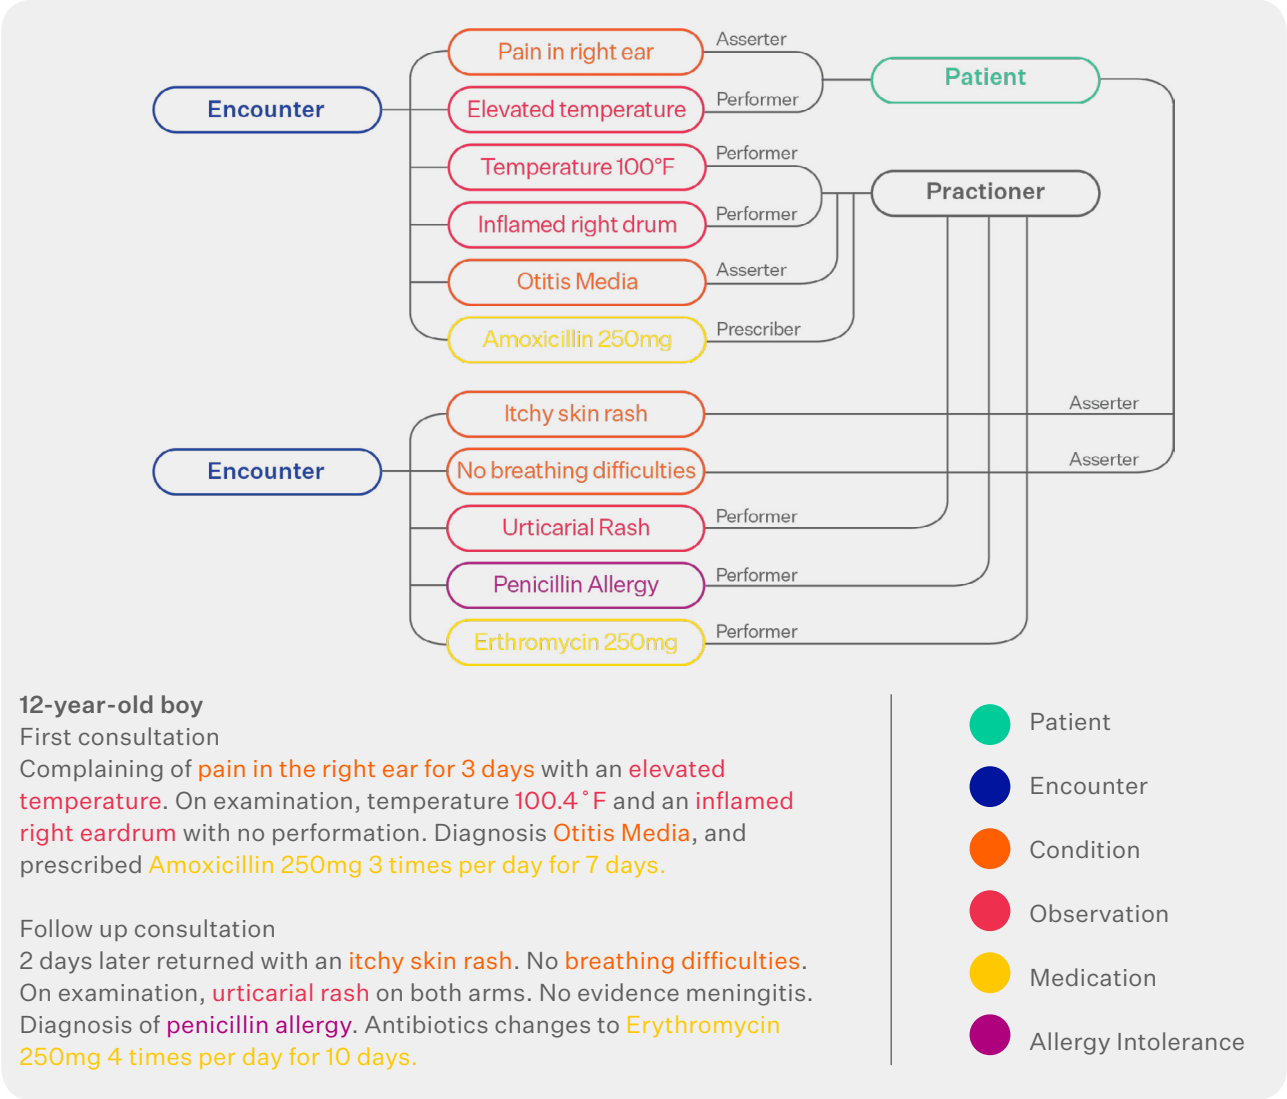

FIGURE 1: Viewing the consultation from the perspective of the web of resources

Like an atom, a resource is itself comprised of sub-parts. Indeed, you can consider that a resource has the following components:

- A textual representation intended for display to a human.
- A fixed set of properties that represent the structured information in the resource – intended for computers to understand. Each property is of a particular 'data type', itself potentially comprised of sub-parts, and it's important to be familiar with them.
  - Examples of these data types are:
    - String – a piece of text.
    - Other simple data types such as numbers, dates and quantities.
    - CodeableConcept – coded information from terminology like SNOMED.
    - Identifiers – such as a medical record number for a patient or an encounter ID for an admission.
- Extensions – which are properties that can be added to a resource to meet a particular need. (These are an important aspect of profiling).
- 'Metadata' – information about the resource such as the date it was last updated, privacy and security tags, and versioning information.

When exchanging information between systems, we think of 4 different 'paradigms' or patterns of use:

1. **Documents** are very familiar in healthcare, with examples such as a Discharge Summary, a Referral letter or a Progress note. Documents are a 'summary at a point in time' about a patient, and are intended to be a part of the clinical record and retained indefinitely.

2. **Messages** are used when one system needs to update another or request that it do something such as ordering a lab test or communicating back the result of that test. Once the target system (such as the Lab order system or the clinical record) is updated, then the message itself can be discarded, other than for audit.

3. **Real-time APIs** (Application Programming Interfaces) allow a system to expose functionality and data that can be used by another application such as a mobile device. For example, a tablet-based application might allow a clinician to access the clinical record and then update it. These APIs (sometimes called RESTful APIs) are becoming increasing important in healthcare delivery, especially with the move towards a more interconnected 'ecosystem' ofdevices and applications. APIs are one of the principal drivers for the development of FHIR.

4. **Operations** extend the API concept by allowing a system to expose complex logic behind simple interfaces. For example, a prescribing service might apply Decision Support logic such as drug interaction or allergy checking before completing a medication order.

All of these paradigms use exactly the same resources internally. This is new to healthcare standards and one of the key benefits of FHIR over the preceding standards. For example it is straightforward to assemble a Discharge Summary (Document) with data collected directly from a Clinician (API) and a Laboratory (Message) and assisted by a context aware application (Service) that assembles the Document.

The FHIR specification defines a fixed set of resource types and their base properties. But there's a conundrum faced by all Standards designers, which is how to decide what properties should be in each resource. This is particularly acute for a standard that is intended to be international in scope, as each country or specialty has a slightly different set of properties that it needs to express. For example a Primary Care Physician describing a diagnosis of Pharyngitis has very different needs to an Oncologist describing a tumor, yet we wish to use the same resource type (Condition) for each one.

Previous standards have approached this by including all the possible properties that anyone anywhere would want to use in their information models. This has resulted in bloated models that are large, complex and difficult to use. FHIR takes the opposite approach by only including those properties used by most existing implementations, and providing a built-in mechanism for safely adapting resources for a particular context of use - a process called Profiling - and it is here that the Clinician comes into their own.

FHIR will allow clinicians to become involved in the delivery of health related IT projects – it is widely recognised that clinician lead IT projects have a higher success rate than technically lead ones. The health IT domain is complicated and FHIR helps to reduce the complexity.

So Profiling adapts a resource for use in a specific context. Some of the things a profile can do include:

- Add a new property (the Extension).
- Change the number of times that an existing property can occur (e.g. that a patient must have one and only one name).
- Indicating that a specific property is not supported e.g. a patient's photo.
- Specify that a coded property must use a code from a specific terminology e.g. LOINC or SNOMED.

That last point, specifying the terminology to be used for a coded property deserves more discussion.

A Terminology (also called a Code System in FHIR) describes a set of concepts with a coherent meaning. A Code System can be small (the possible status's of a lab order) or very large (all possible diagnoses). They can be a simple list of defined 'things' (LOINC test codes) or a complex 'ontology' of inter-related concepts – like SNOMED.

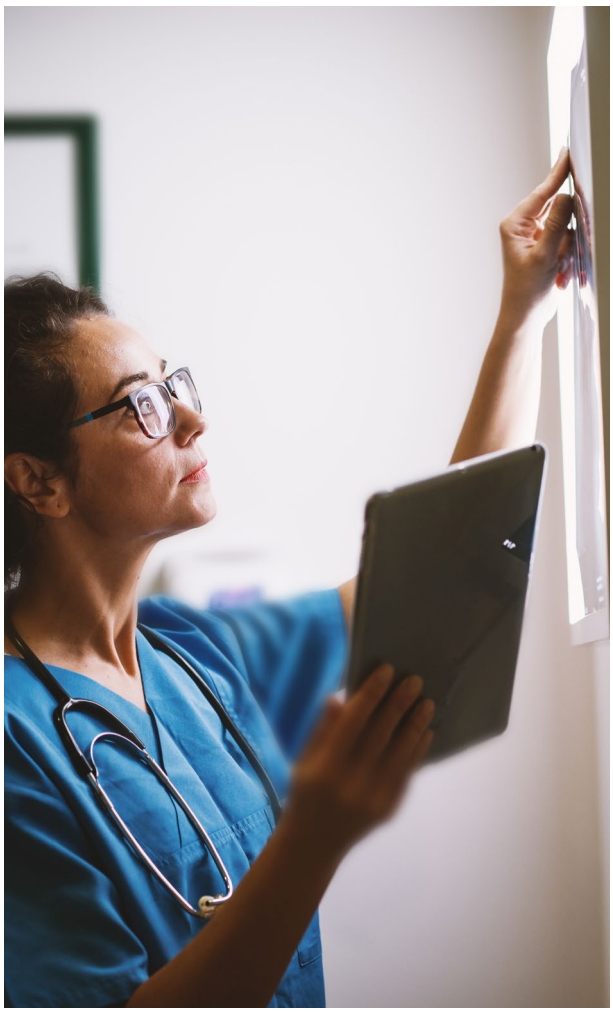

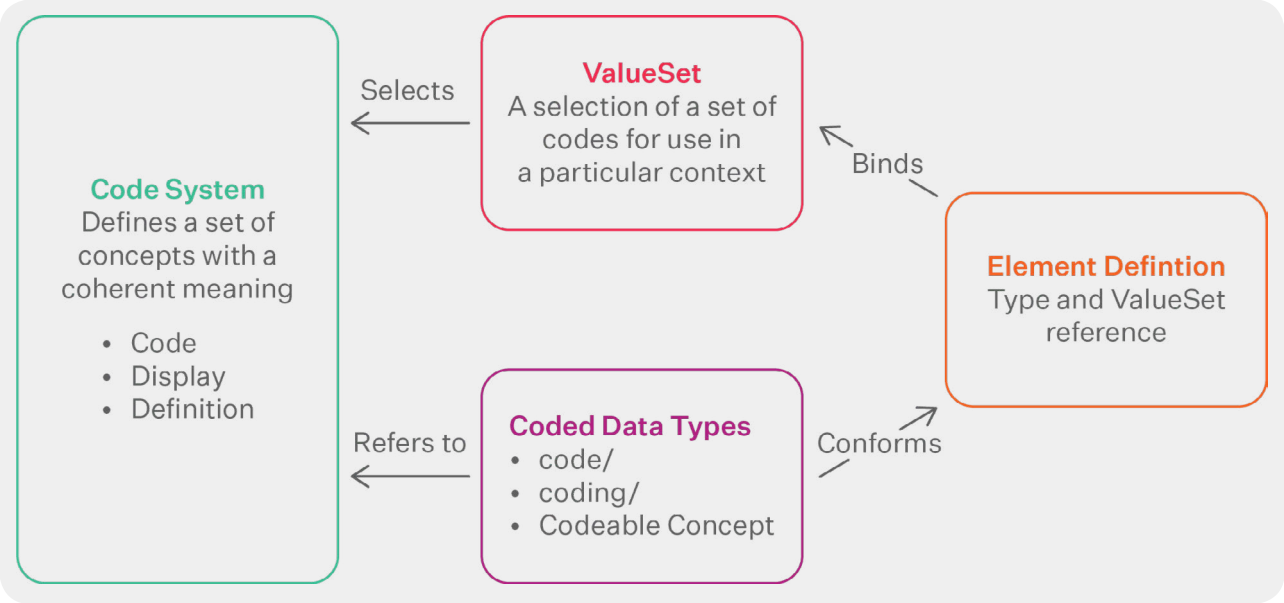

FIGURE 2: Value Set

It is the ValueSet that describes that sub-set of values. Above is a diagram of how the various parts interrelate.

When a coded property is used in a particular context, it is possible to specify the 'sub set' of possible values that make sense in that context. For example, the most common diagnoses seen in the Emergency Room is different to those in an Oncology Department, yet the underlying Code System (commonly SNOMED) might be the same. The ValueSet is used to describe the specific subset.

We use the term 'binding' to describe how the ValueSet defines the set of possible values for a coded property, and 'binding strength' to indicate if a given instance of a resource is allowed to use a value that is not in that ValueSet. In most cases, the ValueSet will contain the most commonly used terms in that context, rather like a 'hint' that makes it simpler to select a concept in a consistent way, but some other term can be selected if needed. However, it is possible to require that an instance value must be from the ValueSet, though this is most commonly used for 'structural' properties like status values, or where for maximum interoperability the set of values should be fixed – such as Allergy criticality.

Profiling Process

When thinking about a specific Interoperability project, perhaps the specification of a pregnancy report, or the format for reporting a suspected drug reaction, there are a number of 'outputs' that will eventually be produced. There will include:

- One or more profiles on resource types.
- ValueSets for the coded properties.
- Examples.
- An Implementation Guide.

How do you start from a clinical need and arrive at these outputs? And in particular, how do you engage with the wider clinical community who do not have an interest in the underlying technology?

There are a number of ways of approaching this problem, one possibility is to use the FHIR Logical model. The Logical Model serves a number of purposes in FHIR, and one of them is as the 'bridge' between the clinical and the IT worlds. It is similar to a resource type, in that it is comprised of properties that use the same data types that resources do, but unlike a resource type the actual properties are not defined by the FHIR specification. They can be whatever is needed to express the information that is to be shared.

The idea is that expert users work with the clinical community to develop the Logical Model unconstrained by the structure of the resource types and profiles that will eventually represent that information in FHIR (and will be understandable by any system that understands FHIR). Once the Logical Model is complete, then specialists in FHIR can use it to build the profiles and other artifacts.

Tooling has started to be developed to assist this process, an example is clinFHIR (clinfhir.com). clinFHIR was developed as a range of open source tools, to be used to educate and assist developers and clinicians involved in the development of the FHIR standard. It serves 2 main purposes:

- As a training tool to help people wanting to learn more about FHIR, visualising how parts combine to represent clinical information in a structured and coded manner. As a development tool with features to build required artifacts.
- The first step is to document the business requirements - what are you wanting to do. There will be aspects of workflow and/or process involved in this step, but the parts we are focusing on here is the information that needs to be shared. There could potentially be more than one model. Domain knowledge is required for this step.

The next step is to use the logical modeler in clinFHIR to create a structured representation of the information to share – the model. Familiarity with the FHIR data types, and ideally the basics of FHIR (especially referencing between common models/resources) is preferable but really domain knowledge is the most important.

You need to decide if you are modeling what will become a single resource (e.g. an allergy intolerance) or a collection of resources – like a list, document or message. In the latter case you might want to split the design up into a number of referenced models to make the next stage simpler.

This step can take some time as getting agreement from multiple people is not always straightforward. There are a number of outputs from this stage:

- The models themselves will become profiles – both constraints and extensions on resource types.
- The ValueSets that will be needed for the coded elements.

The final stage is to create or locate the FHIR artifacts – ValueSets, Code System, Profiles and Extension Definitions that represent the models in FHIR-space. This will need to be done by those with a good understanding of FHIR, and will be 'wrapped' up in an Implementation Guide that will include other artifacts such as scenarios, examples and more detailed documentation.

Over the page is an example of a logical model - an Adverse Drug Reaction report, and would likely result in a number of profiles:

- The overall model will be a profile on AllergyIntolerance.
- The 'drug' property would be a profile on Medication.
- 'Reporter' would be a Practitioner, and 'patient' a Patient.

To summarise, the parts of FHIR that a Clinician involved in an interoperability project should be aware of (and contribute to) include:

- Resources - the basic information units in FHIR.
- References between resources – telling the clinical story.
- Properties and Data types – the internal parts of the resource.
- Profiling – making FHIR work in different clinical contexts.
- ValueSet – describing the set of concepts for a coded property in a specific context of use.

Conclusion

FHIR is built around units called "resources", the basic unit of interoperability. Examples include a Patient, Condition, or a Procedure. FHIR covers both the format of information and how data is exchanged, so it is both a Model and an API. It promises to make health information easily and securely accessed from any device, anywhere. FHIR will help to break down information silos that exist in healthcare. FHIR could support access to vital data at the push of a mobile app button. It will even support the creation of an "app store" of independently developed mobile applications, because it supports the collection of data and availability via standardised APIs.Clinicians and patients will benefit from the FHIR standard as it improves access to a more complete, higher quality electronic health care record, by being able to include data from traditional sources like laboratory results, and evolving sources like genomic information.

FHIR promises to revolutionise sharing of healthcare information. The health IT domain is complicated and FHIR helps to reduce the complexity. It is now time for clinicians to become more involved in FHIR related IT projects. In this paper we have discussed the fundamentals of FHIR, the process involved with developing a FHIR interoperability project and artifacts. The clinFHIR tool can be used to educate clinicians about FHIR and enable them to be actively involved and blaze through FHIR healthcare IT projects.

Dr. David Hay - Biography

David is a Product Strategist for Orion Health. He is also active in the international standards community as the chair emeritus of HL7® New Zealand and is a co-chair of the FHIR® Management Group, charged with guidance and development of the latest HL7 standard. David graduated from medical school, then moved into the Health IT sector. He started a company developing Practice Management Software to GP's (GPDat). This was the first such program in New Zealand to receive electronic laboratory results.

Leaving the vendor space, David worked at EDS as a Solutions Architect, before returning to health IT working as an Architect (Solutions and Enterprise) for the Auckland based services management organisation 'healthAlliance' producing a number of innovative solutions: an internal eReferrals application, a community based Case Management solution, plus an application to track and report on clinical tasks within the hospital. All of these applications remain in current use today. While there, he also participated in a number of national programs, including the medical records transfer project GP2GP, ePrescribing, and was one of the authors of the Interoperability Reference Architecture that governs information flow within the health sector.

He currently serves on the Health Information Standards Organisation (HISO) committee, which provides technical advice on standards to the New Zealand Digital Advisory Board. He has provided outstanding leadership towards helping to create and evangelise the innovative FHIR standard, and frequently writes about FHIR on his blog fhirblog.com and has developed a range of open source tools widely used to educate and assist the developers of FHIR clinfhir.com.

David is a true champion of the New Zealand IT healthcare sector, he has spent many years both professionally and personally designing, educating, and advising on health informatics. In recognition of this David has been awarded the Excellence in Health Informatics award 2016 by ITx New Zealand.

His ongoing commitment to FHIR has helped to provide the momentum needed for international adoption. FHIR represents a major standards upgrade that will boost access to health information, which will improve the access to a patients health information globally.

ORION HEALTH