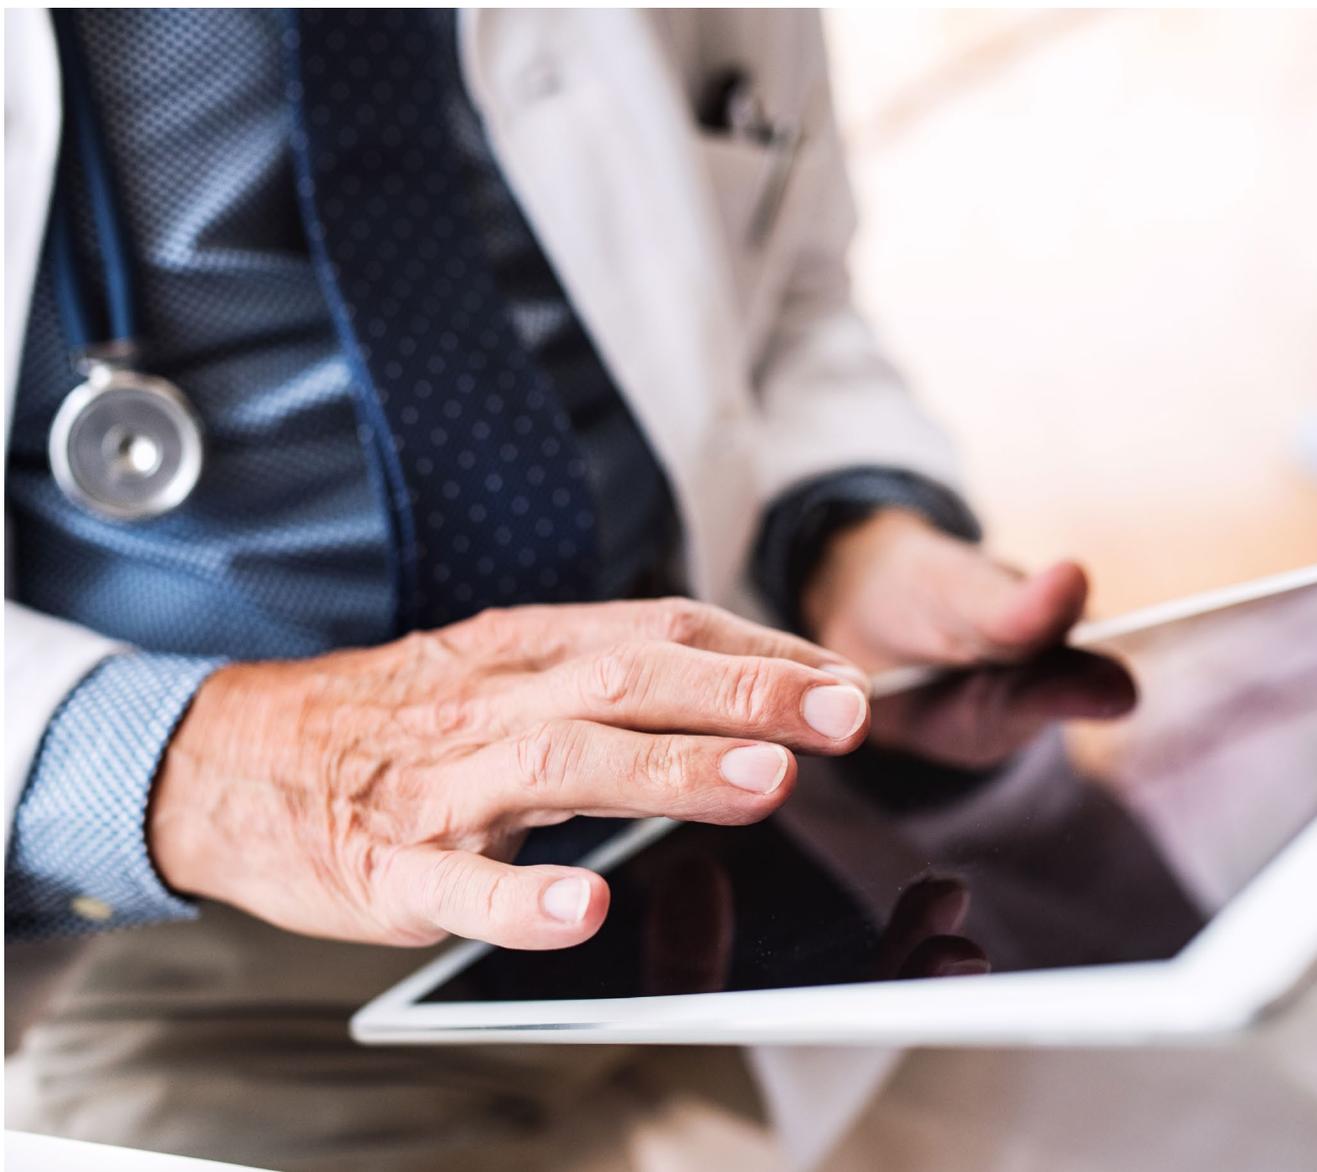# FHIR – Sparking Innovation in Health Information Sharing

Orion Health White Paper
Dr. David Hay
Product Strategist
112015

ORION
HEALTH

How FHIR Benefits Clinicians and Healthcare Executives

Health information sharing is becoming increasingly important, as all those involved in the delivery of care to patients now expect that the information they need to be available to them at the point of care and not locked into the 'silos of collection' which has historically been the case. At the same time, we are seeing a huge increase in the amount and source of that data as genomics and personal health collection devices become affordable and more prevalent. And this data is not only required for direct care delivery, as the value and need for population-based analytics is also increasing.

In response to these needs, a new healthcare standard has emerged from HL7. Called HL7® FHIR® (Fast Healthcare Interoperability Resources) and pronounced "fire"), it is continuing to gain momentum and acclaim as those who use it realize the elegance of its design and ease of use. At this stage, most of the people involved in its development, which is a community-oriented process, have been technical implementers and those clinicians already involved in healthcare informatics. As the standard continues to mature and especially as tooling becomes available, it has become feasible for other clinicians, motivated patients and healthcare executives to become more involved in the on-going development.

This paper examines how FHIR is sparking innovation in health information sharing and how it can benefit clinicians, healthcare executives and informed patients. The paper starts with the main components of FHIR, then moves on to profiling, which allows the standard to be adapted for real-world scenarios. The paper will finish with a discussion of how different stakeholders can benefit from the FHIR phenomenon.

What is FHIR?

There are two main aspects to consider:

- The way in which information is represented, known as "resources".

- The manner in which it is shared, whether in a document like a Discharge Summary, in a message between systems in a facility or via a real-time API (Application Programming Interface)-based exchange.

The FHIR resource

A resource is the basic unit of information exchange, the smallest thing that makes sense to be shared. Resources are deliberately kept simple by including in the core resource only the data commonly exchanged by systems today, but FHIR provides a built-in extensibility mechanism by which the needs of real-world scenarios and specialties can be met.

FHIR resources cover the full range of healthcare interoperability requirements, including:

- Clinical data such as problems, medications and allergies.

- Administrative information such as patients, appointments, encounters and financial resources.

- Infrastructure resources such as lists, orders and devices.

There is also active work within the community to integrate the large volume of genetic information required to support the new "precision medicine" that offers so much promise in targeting treatment, especially for individuals with chronic diseases or malignancies.

```xml
<Patient xmlns="http://h17.org/fhir">
    <id value="100"/>
    <meta>
        <lastUpdated value="2014-11-13T11:41:00+11:00"/>
    </meta>
    <text>
        <status value="generated"/>
        <div xmlns="http://www.w3.org/1999/xhtml">
            <p>Henry Levin the 7th</p>
            <p>MRN: 123456. Male, 24-Sept 1932</p>
        </div>
    </text>
    <extension url="http://orionhealth.com/fhir/StructureDefinition/trials">
        <valueCode value="renal"/>
    </extension>
    <identifier>
        <use value="usual"/>
        <label value="MRN"/>
        <system value="http://www.goodhealth.org/identifiers"/>
        <value value="123456"/>
    </identifier>
    <name>
        <family value="Levin"/>
        <given value="Henry"/>
        <suffix value="The 7th"/>
    </name>
    <gender value="male"/>
    <birthDate value="1932-09-24"/>
    <careProvider>
        <reference value="Organization/2"/>
        <display value="Good Health Clinic"/>
    </careProvider>
    <active value="true"/>
</Patient>
```

Resource Identity and Metadata

Human-readable Summary

Extension with URL to definition

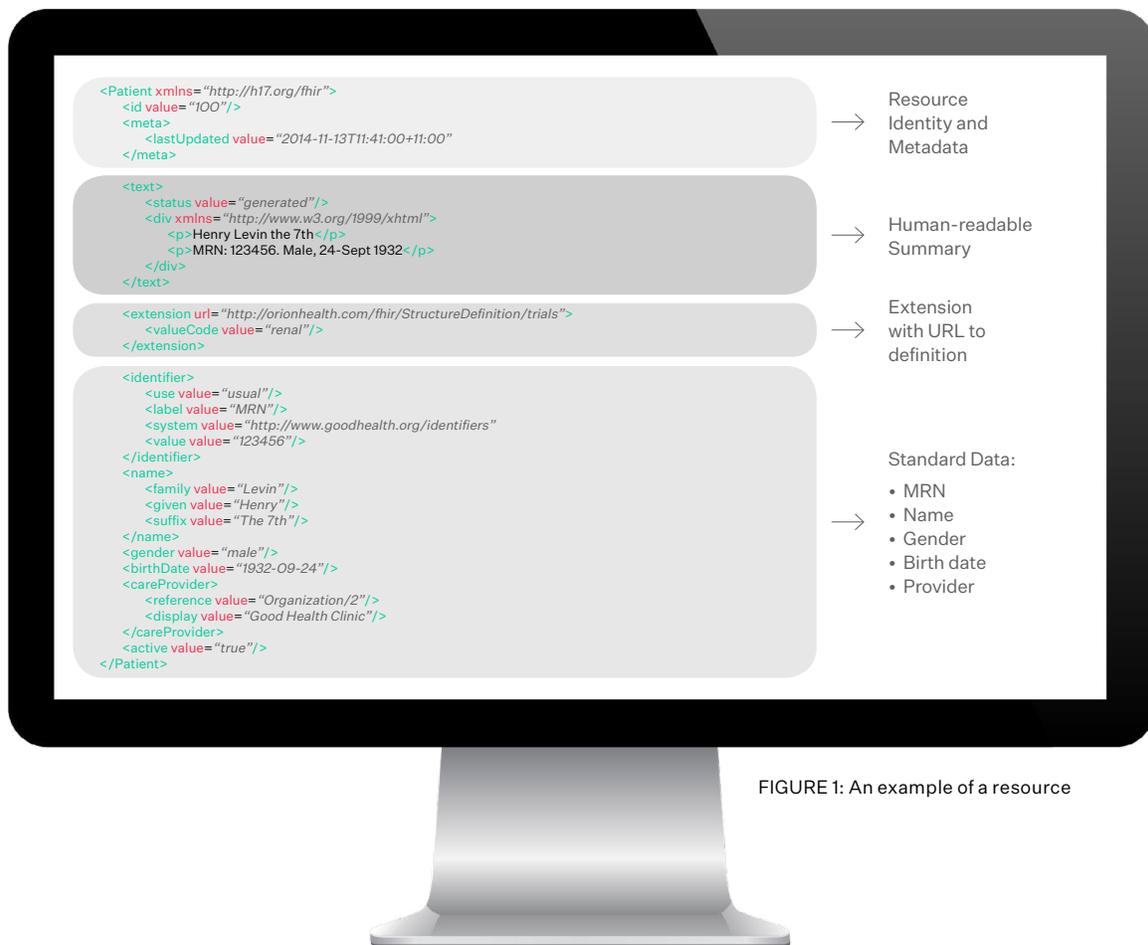Standard Data:
• MRN
• Name
• Gender
• Birth date
• Provider

FIGURE 1: An example of a resource

This example is a Patient resource, and it shows the four common parts present in each resource:

- The metadata about the resource, when it was last updated, the ID etc.

- A human-readable summary. All resources can represent their content in a form that can be displayed to a human.

- Extensions, which are optional "additional" properties beyond that defined in the core resource definition. Extensions must contain a reference to their definition. This is important, as it means that any recipient can go and find out what it means if they are not familiar with it.

- The structured elements, as defined in the specification.

Although not intended for human consumption, the simplicity of the structure brings great benefit to implementers as it lowers the bar for understanding how complex healthcare information is represented, and makes it easier for people without a great deal of exposure to healthcare information to rapidly become productive in FHIR implementations.

A web of resources

To represent real clinical data, you need to have more than one resource. FHIR supports this with references between resources. In the example below, the details of a procedure are shown. We can determine:

- What the procedure was – e.g., an Appendectomy (Procedure resource)

- On whom it was performed (Patient resource)

- When and where it was done (Encounter)

- Who performed it (Practitioner)
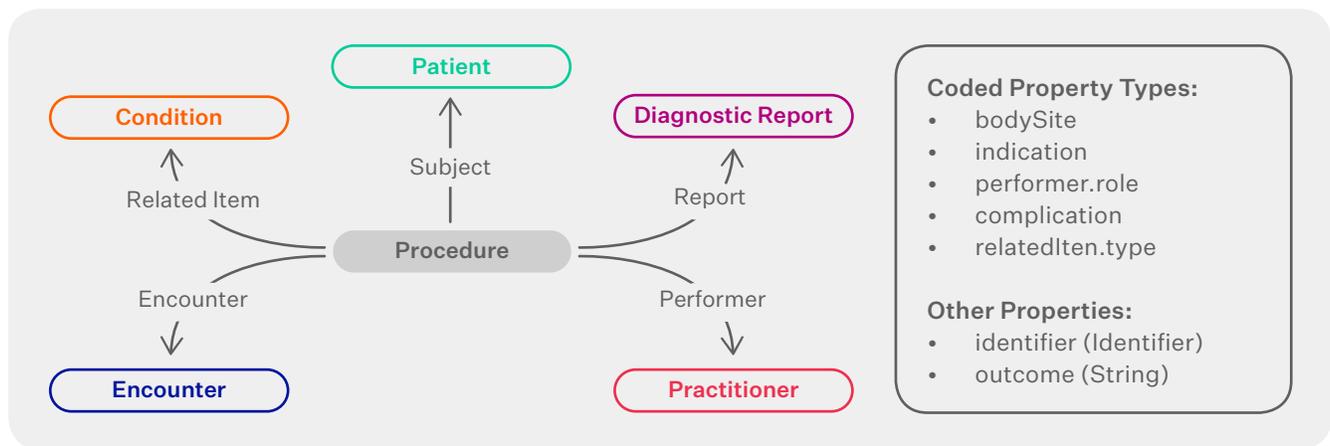
- A pathology report on the removed tissue (DiagnosticReport)

FIGURE 2: An example of a procedure

This ability to reference between resources means that you can use FHIR to describe any healthcare event in a way that a recipient can easily interpret.

For a more complete example, here is a textual rendition of a simple clinical scenario, with the FHIR resources linked to the elements of the consultation:

12-year-old boy
First consultation
Complaining of pain in the right ear for 3 days with an elevated temperature. On examination, temperature 100.4˚F and an inflamed right eardrum with no performation. Diagnosis Otitis Media, and prescribed Amoxicillin 250mg 3 times per day for 7 days.

Follow up consultation
2 days later returned with an itchy skin rash. No breathing difficulties. On examination, urticarial rash on both arms. No evidence meningitis. Diagnosis of penicillin allergy. Antibiotics changes to Erythromycin 250mg 4 times per day for 10 days.

Patient
Encounter
Condition
Observation
Medication
Allergy Intolerance

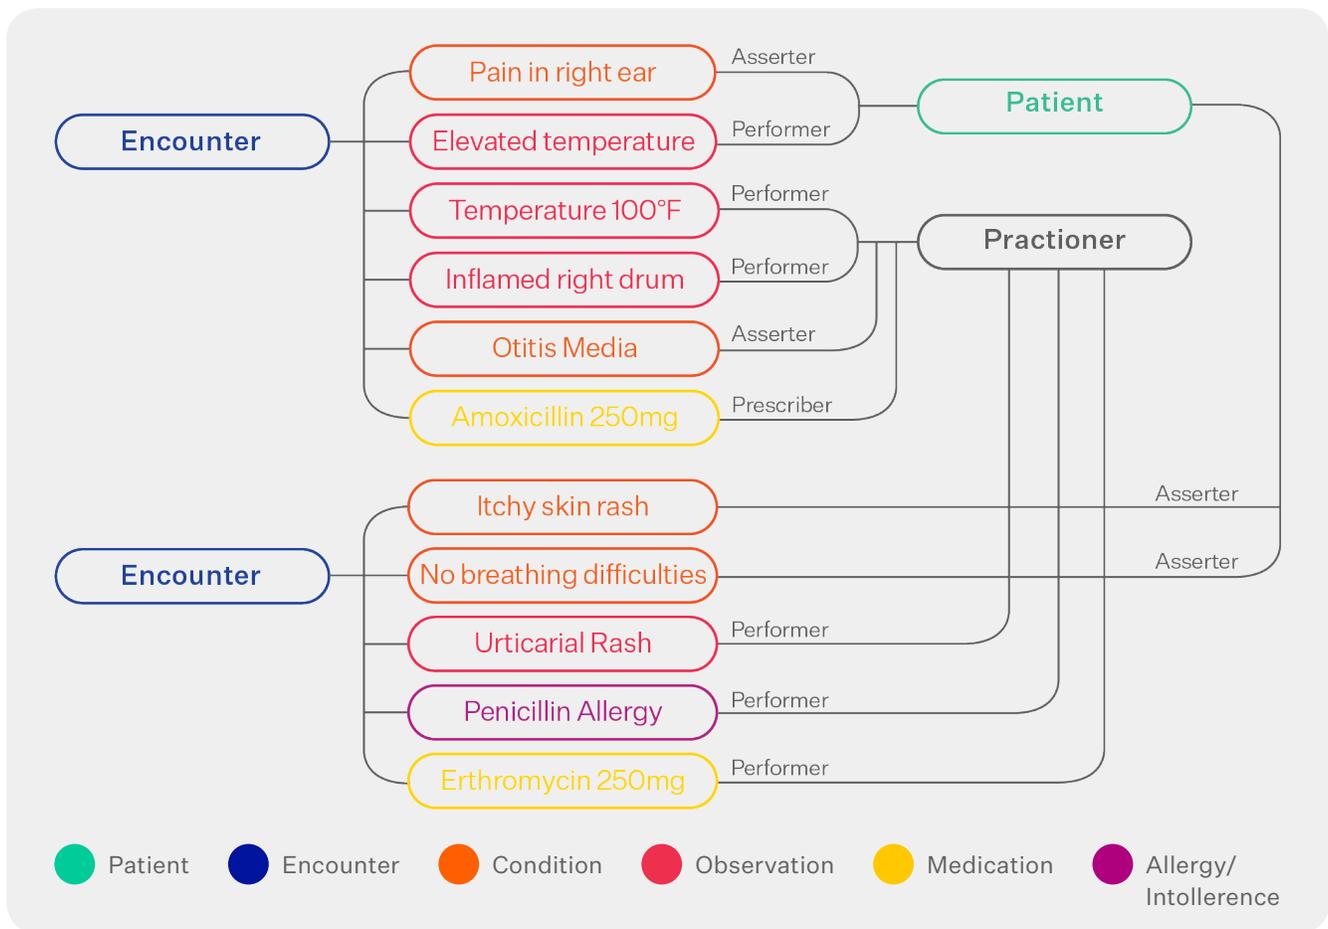FIGURE 3: A textual rendition of a simple clinical scenario

FIGURE 4: Viewing the consultation from the perspective of the web of resources

Data Types and Terminology

The elements within the resource, whether the defined structured data or extensions, each have a specific data type, or "kind" of information. There are, broadly, three different categories of data type.

- Values in the Resource: These can be simple values like a string, a number or more complex values like an Address or a Name.

- Resource References: As in the example above, these are how individual resources can refer to other ones to create the web of information.

- Coded values: These are elements where the actual value is drawn from a set of possible values, often described as a Terminology.

Terminologies are extremely important in healthcare, as they define the "concepts" that are being exchanged; i.e., the procedure being performed in the example above, a specific diagnosis or a medication.

There are a number of terminologies available for specific purposes, and FHIR can support any of them. SNOMED is commonly used for a diagnosis or problem, LOINC for lab results and RxNorm for medications.

A coded datatype in FHIR (like Procedure.code) is bound to a set of values from a particular terminology using a ValueSet. The specification defines some of these bindings, but in most cases, these can be changed during Profiling, as they are usually specific to a particular context of use.

Profiling

Profiling in FHIR is one of the most important concepts for a user to understand. By design, the resources are kept small, and contain only the most commonly used elements for a given resource. For example, a patient resource has the patient's name, address, date of birth and gender. However, a given implementation might also need to include religion or race, neither of which are included in the base resource. There might also be elements defined in a resource that are not needed in a given implementation, such as Patient.animal or Patient.photo.

It is the profile that takes the base resource, and adapts it to the specific requirements of an implementation, and it does so in a way that is computable, which means that it is able to be understood by a computer. This supports features like being able to validate that a given resource is conformant to a profile or to automatically generate a user interface from a profile.

Therefore profiles allow a user to take the base resources and make them perfectly fit their needs in a way that others can understand. Examples of what a profile can do include:

• Add a new element to a resource (such as Patient.religion)

• Remove an unsupported element from a resource. There are some limitations – you cannot remove an element that is required in the base resource, for example.

• Change the multiplicity of an element (again, with some restrictions). For example, Patient.name allows more than one name to be present. A profile can limit this to a single name only, and can make the name required.

• Change the terminology binding to a

different ValueSet. This is really useful in specialist applications. For example, if you were building an application for use in an Emergency Department, you might want to create a subset of SNOMED (using a FHIR ValueSet) that contains the most commonly used diagnoses. The User Interface can then present this list by default, making it simpler for the clinician to select the diagnosis, and improve consistency between implementations of that application.

Because FHIR has been built to allow controlled customization via profiling, it is likely that a recipient will receive a resource that has been profiled in some way. How can they understand what those adaptions actually mean? The answer is that FHIR defines "'Registries" on the Internet where these definitions (which are FHIR resources themselves) can be stored for a consumer to access if needed. Anyone can establish a registry (HL7 will be providing one for implementers to use), and the specification allows for these to be distributed across different countries for resiliency and performance.

Sharing resources

There are a number of different ways in which resources are shared (commonly called paradigms of exchange). These are:

• Messaging, where one application needs to inform another of some data, or instruct it to perform some task. Once the client has acted on the message, it can be discarded. An example of this is the ADT messaging in hospitals where the hospital administration system notifies the Laboratory system of a patient's admission.

- Documents, such as a Discharge Summary. These are summaries of clinical information about a patient at a point in time that need to be retained indefinitely.

- Service APIs (Application Program Interface) which support real-time access to data. These are relatively new in the healthcare world, and until FHIR, have not been readily available.

FHIR supports all these paradigms, and because the resources are the same in each one, it is easy to acquire information in one paradigm and share it in another. For example, a clinician might update a Problem List in real time through their mobile device, which is subsequently incorporated into the Discharge Summary document.

The ecosystem

The main parts of FHIR have been described, so how do they fit together to improve health information sharing? How does this translate into the bigger picture and the health information ecosystem? What do we mean by "ecosystem" and why would we want one?

Based on the natural world, an ecosystem is all about how multiple independent "things" work together in a common community. Applied to healthcare information technology these independent parts of a whole are represented by the individual applications, this means the individual applications that clinicians and patients use to record and share that data with others.

An ecosystem is the opposite of a single big system – one common application that everyone uses. While the latter is simpler to implement, and appropriate in some settings, it does mean that it becomes hard to do anything different; anything that the application has not been built to do. The ecosystem, on the other hand fosters

innovative applications (especially mobile applications) by exposing information and services so that the data can be repurposed. Examples of these kind of applications include:

- A diabetes application that allows the patient to enter their daily blood sugar levels and medication. This information can be shared with their primary care physician or diabetes specialist.

- An application for a clinician that allows notes to be reviewed and Orders submitted.

Benefits to users

What are the benefits to clinicians and their patients?

- Because it is easier and safer to exchange and collect high quality information, clinicians will have a more complete clinical picture when managing patients. This will have benefits for both the clinician and the patient in saved time and improved quality of care.

- The increasing range of applications that can securely access information will make it easier to develop "purpose-specific" applications. The information needs of, say, a surgeon, are quite different to the needs of an anaesthetist. While there are many such apps available today, in most cases, their access to information is limited, often only to the data entered through that application. This greatly reduces the potential value to the end user. Being able to use FHIR to access a common pool of data will greatly enhance their value. To further facilitate this, the SMART standard extends FHIR and security standards to do this in a common way. See smarthealthit.org/smart-on-fhir/ for details.

- This range of applications will also benefit the patient, making it simpler for them to participate in their own healthcare as an equal partner. Because mobile applications will be crucial in servicing this sector, FHIR's excellent fit in that space will further improve the ability for innovative vendors to create purpose-fit applications.

- Validation, and therefore, information quality, becomes simpler as all resources, even profiled ones, can be automatically verified as being "conformant". The FHIR project has made sample code available that can do this, and this will improve the overall quality of exchange.

- Tooling is being developed that puts a user-friendly face on top of FHIR. Reference code, libraries and test servers have always been a part of the FHIR project. These are reasons why developers and implementers have become so enthusiastic about FHIR, and now there is tooling aimed at the end user. The Netherlands-based company Furore are producing a number of tools for use with FHIR including a profile builder (they can be found at fhir.furore.com/) and there are also on-line tools like clinfhir.com with a more educational focus.

Why do innovators and implementers embrace FHIR?

FHIR makes it easier for implementers by using familiar tooling and an understandable format that doesn't require domain knowledge, plus an API style extensively used internationally. From the perspective of interoperability, a fixed set of resources (that can be extended for specific scenarios) offers a number of advantages.

- By standardizing the form and content of the resource, variability from source systems is reduced, thereby reducing the mapping (and time) required to integrate. A consistent and understandable format makes it easier for implementers to perform the mapping that is required. While not a design goal for FHIR, having wire representations that humans can understand reduces the possibility for error and the amount of testing required.

- The availability of predictable interfaces, both in terms of content and APIs make it easier to develop applications, as much of the design work has been done. This leaves innovators and implementers free to focus on the value add of their applications.

- Using standards-based technologies such as XML, JSON, REST and Oauth lowers the barrier to entry for new participants.

- The specification is publically available on the Internet and free for all to use. It is fully hyperlinked with hundreds of examples. The goal is for each property of each resource to have at least one example that demonstrates its use.

- There are a number of open source libraries and online test servers developed by the community that further reduce the cost of entry. These libraries are being used by many of the established vendors as they provide FHIR interfaces for their products.

- There is a large support community that is available through diverse channels including Skype, email, list servers and "Stack Overflow," with more channels planned. It is not unusual to get answers within minutes of asking.

- Face-to-face meetings (often called Connectathons) are regularly held. There are meetings for both implementers and

clinicians, and have a dual purpose:

- To ensure that FHIR is fit for purpose by testing proposed changes by implementers before the specification is updated.

- As a networking and training venue for those new to FHIR. It is not unusual for a participant completely new to the specification to develop working software in less than a weekend.

- Many vendors are starting to implement FHIR, as this increases the potential market for their applications. The SMART standard uses FHIR and OAuth2 to provide secure access to data held in EMR systems.

- The conformance resources that FHIR defines allow a server to "advertise" the resources and profiles they support, making it easier to match an application with a server.

Improving the viability of FHIR-based applications encourages innovation in the healthcare information arena, increases competition and ultimately, will help to reduce overall costs involved with health information sharing.

What are the benefits to Healthcare Organizations?

The wide acceptability of FHIR will result in faster implementations of solutions that are cost effective to maintain. This will happen whether or not the implementers are external vendors or internal development capability. The standards-based interfaces will further reduce the risk of vendor lock-in, making it easier to replace a non-performing vendor with an alternative. Data will no longer be trapped trapped in silos which can only be accessed by the host applications.

Conclusion

This paper examines how FHIR is sparking innovation in health information sharing, and how it can benefit clinicians, healthcare executives and informed patients. We have covered the main components of FHIR and profiling – which allows the standard to be adapted for real-world scenarios – and how different stakeholders can benefit from FHIR.

In conclusion, FHIR is truly a phenomenon in the way that it has brought together people from many different areas in an active community focused on making healthcare interoperability as easy as it can be. It is now at a level of maturity where clinicians and healthcare executives, who do not normally get involved in standards development, can actively participate.

Further information can be found at
wiki.hl7.org/index.php?title=FHIR

Dr. David Hay — Biography

David is currently a Product Strategist for Orion Health. He is also active in the international standards community as the chair of HL7 New Zealand and is a co-chair of the FHIR Management Group, charged with guiding the development of the latest HL7 Standard. David graduated from medical school in 1981, afterwards moving into the Health IT sector. He currently serves on the Health Information Standards Organization (HISO) committee, which provides technical advice on standards to the National Health IT Board , and frequently writes about FHIR on his blog
at fhirblog.com

ORION
HEALTH