

# FHIR Ignites Healthcare Sharing



**Orion Health White Paper**  
Dr. David Hay  
Product Strategist  
052016



## Promoting interoperability in the healthcare ecosystem with the HL7 FHIR standard

The delivery of quality healthcare in the modern world is absolutely dependent on the availability of quality information. This is true whether the information comes directly from a clinician, monitoring by a Care Coordinator or through an anonymised population analysis. The problem is that data is held in many different places – often only by the system that collected it in the first place – and often the structure and content of that data is focused on the needs and formatting of the collecting system, rather than on formats more suitable for wider sharing.

The ability to exchange information between systems in a timely and understandable manner has always been important, and as the volume and type of health-related information increases, it is becoming even more important and difficult to achieve.

Historically, much of this exchange has been performed using standards developed by HL7®, including version 2 messaging, which has been in use since 1987. More recently, Clinical Document Architecture (CDA) and derivatives such as Consolidated CDA (CCDA) have become widely used – especially in the United States, but also worldwide. However, these standards can be complex to implement effectively, and are outmoded in the real-time connected world required by mobile and personal device applications.

In recent years, a new standard has emerged from HL7. Called HL7® FHIR® Standard, Fast Healthcare Interoperability Resources (FHIR, pronounced “fire”), it has attracted enormous interest from the wider healthcare community at a pace and level of interest unmatched by any previous standard.

This paper examines the drivers behind FHIR, gives an overview of what it is, how to learn more, how it will benefit healthcare information exchange and provides a future perspective on where FHIR is heading.

## Background

There is an increasing need worldwide to make healthcare information available to those who need it to deliver care to patients – including the patient themselves. This is driven by a number of factors including:

- Management of chronic conditions such as diabetes and heart disease requires coordination between multiple groups to achieve the best possible outcomes for patients – and reduce the incidence of costly secondary complications.
- An aging population, with an expectation that they will remain healthy and active as they age.
- An aging workforce, reducing the people available to deliver care and meaning that efficiencies in care delivery has become urgent.
- The increasing cost and options for healthcare delivery, and the increasing demand from patients for these options.

At the same time, the amount – and nature – of clinical information is increasing. Apart from the usual sources, such as improved diagnostic equipment and tests, we are seeing a surge in information from the patients themselves in the form of device data such as wearable devices (often fitness-based) and devices that can electronically upload measurements (weight scales, glucose meters, blood pressure monitors). Combined with the dramatic cost decrease of genomics-based data, the advent of “personalised medicine” seems certain.

All this information needs to be collected, stored, analysed and shared as the diagram indicates:

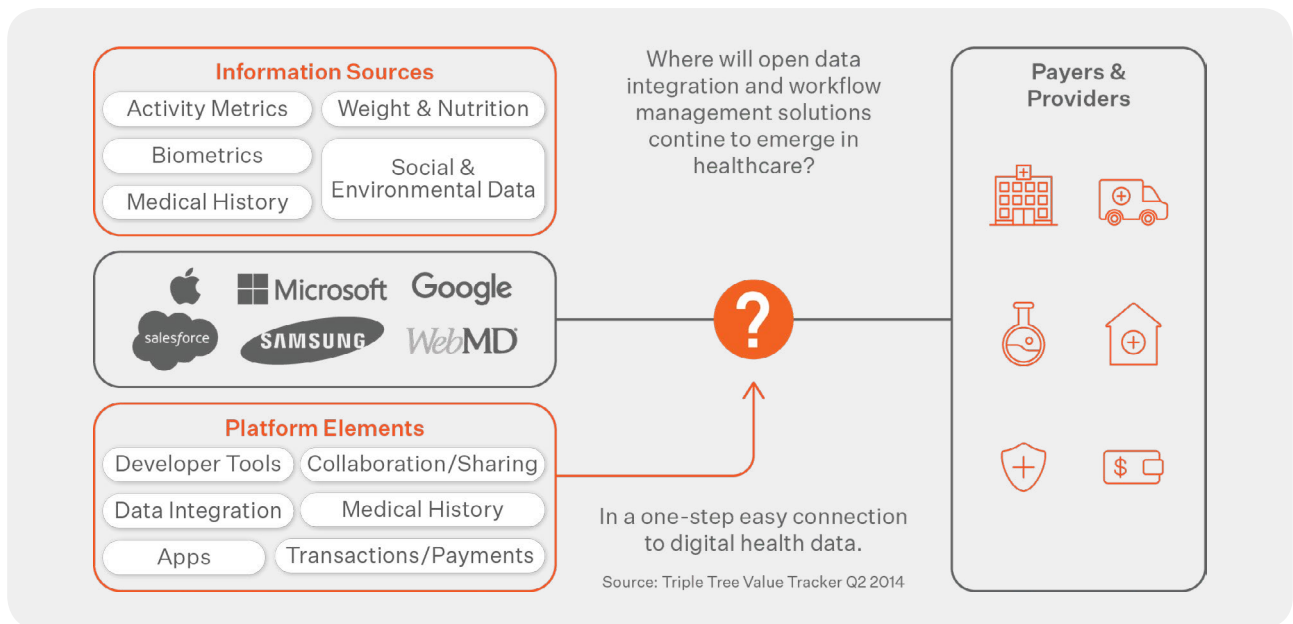


FIGURE 1: Information Processing

## What is FHIR?

Simply put, FHIR is a new interoperability standard from HL7 whose purpose is to make it easier to exchange all healthcare data between systems in a straightforward yet secure manner. It is built around the concept of “Resources” – self-describing, discrete “chunks” of data that “make sense” in the healthcare environment – examples of Resources include Patient, Condition, Procedure, Medication, Allergy, Observation and Appointment.

Each resource has a small set of associated properties (like name, date of birth, gender), but has the built-in ability to be extended and profiled to meet specific use cases. This allows the core resources to remain small and manageable, yet expand to meet any requirement. Further, the specification supports “discoverability” of these extensions so that the recipient of an unknown extension in a resource is able to determine what it means.

Resources are linked together in a network or graph of related resources as the diagram below shows:

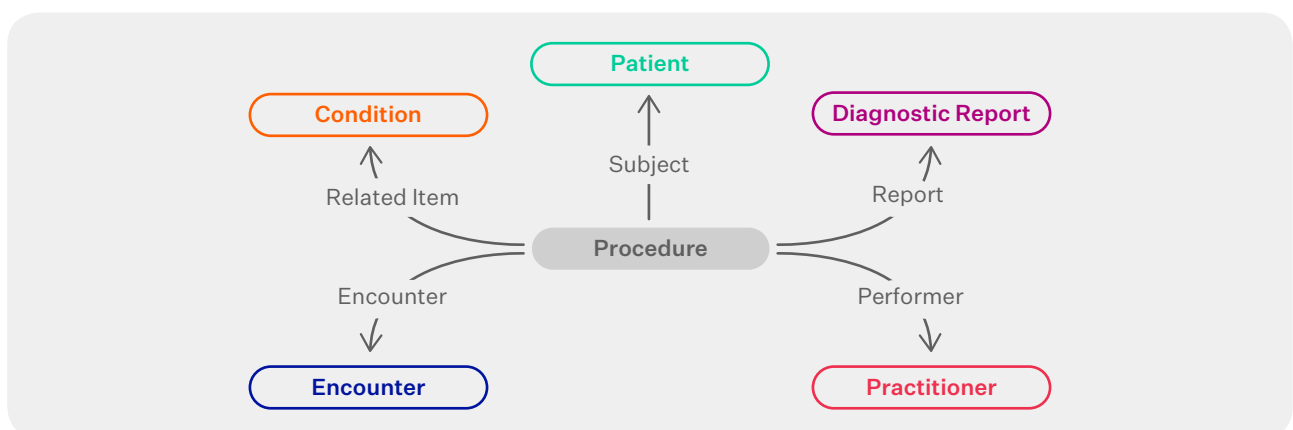


FIGURE 2: Procedure Example

In addition to describing what a resource is (often called the “Model”), FHIR also describes how to locate and exchange them – and the Application Programming Interface (API) that can be used to accomplish this. There are a number of different paradigms of exchange that are supported in FHIR:

- **Lightweight**  
Online real-time exchange, as used by vendors like Google, Twitter and Facebook (often described as RESTful exchanges). This is especially useful for mobile applications, but extensively used by browser applications as well.
- **Documents like Discharge Summaries, Progress Notes and Referrals**  
There is active work within HL7 (and outside it) to ensure that the currently used document standard – CDA – can be represented in FHIR documents. Indeed, project Argonaut was funded by a group of major US vendors and providers with the express purpose of ensuring that the design work required would be completed in a timely manner.
- **Messages like the HL7 v2 ADT messages, which are very widely used within institutions worldwide**  
Messages are often described as “ephemeral” – they are a notification that something has happened (or needs to happen) from one system to another that is discarded once the target system has been updated.

- **More customised services where the target of the service needs to perform more complex processing than simple data exchange**

For example, prescribing services that may involve patient-specific decision support such as drug interaction checking. In all of these paradigms, the same set of resources is used – for example, an Observation resource may be received from a laboratory in a message, and then included in the subsequent discharge summary document.

There are many architectural approaches that can be taken with FHIR – the diagram below shows one example of putting a FHIR interface in front of an existing data source. This example also shows some of the other standards that FHIR leverages.

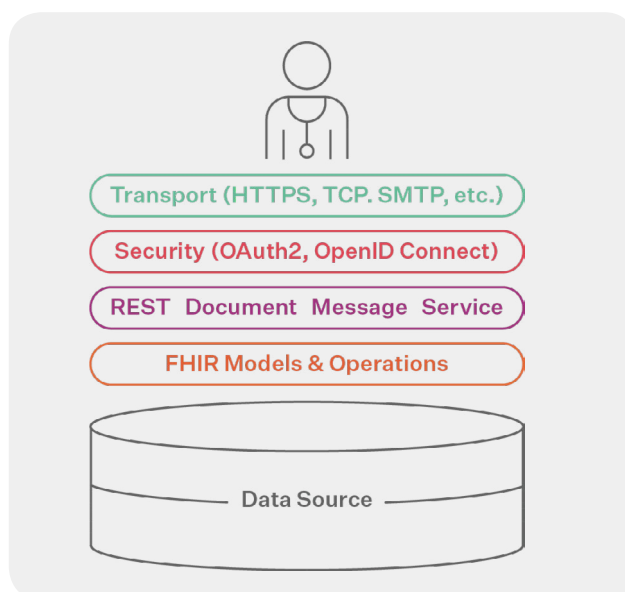


FIGURE 3: FHIR Interface Example

## So why has FHIR gained the attention that it has?

There are three aspects to the answer:

- Implementer focus
- Ease of use
- Community

### Implementer Focus

FHIR is focused firmly on the implementer of systems. For example, in considering the core contents of a resource, this is expressed in the “80%” guideline (and it is only a guideline), where a property resource is added to the core resource only if it is currently being used by 80% of systems – any additional properties would be implemented with extensions. In other words, FHIR aims to support real-world current use and what systems are currently doing, rather than what standards developers think they should be doing. FHIR’s focus on the core resource is possible because of its ability to extend resources to meet specific use cases through the extension discovery and profiling mechanisms.

Indeed, FHIR has involved implementers from the very beginning using events such as Connectathons, where the developing standard is tested by the implementers who will be using the standard, with training supplied to them and their feedback incorporated into the standard.

By supplying open source libraries that developers can use, and the provision of Internet-available test servers, FHIR is available for anyone to learn.

### Ease of use

One of the difficulties that a new implementer faces is the need to understand the complex domain that is healthcare interoperability. Often, significant study must be done before someone can become proficient using healthcare standards, as

much of the language and tooling is specific to the healthcare domain in which it is used. By leveraging standards and tooling that are used outside of the healthcare domain and are therefore familiar across industries, new implementers can build on existing skills rather than needing to develop entirely new ones specific to healthcare. It is not uncommon for people to be productive within days of starting to develop in FHIR, rather than weeks or months. This makes FHIR a cost-effective option in terms of an implementer’s start-up time.

In addition, by keeping the core resources simple but also providing a simple discoverable extensibility mechanism, implementers can “start small,” yet have the confidence that FHIR can grow with their needs.

### Community

From the beginning, FHIR has always been a community-developed product. While this is true of many previous standards, the FHIR community has been able to leverage the social tools that have become available in recent years – including Skype, list servers, blogs and the open nature of the Internet itself. The standard is freely available online to anyone who wishes to view it, and the ability to make comments is both welcomed by the community and supported by the specification – for example, every page has a link that allows anyone to make a comment, and these comments are managed within an openly available source control system.

This open community ethos has also meant that other Standards Developing Organisations (SDO) have added their expertise to FHIR development. IHE (Integrating the Healthcare Environment) and openEHR have both contributed to FHIR, as have others, such as the SMART project (<http://smartplatforms.org/smart-on-fhir/>).

Another example of how the FHIR community is growing is “project Argonaut” – a vendor- and provider-sponsored project created to assist in the development of FHIR – especially in the alignment of FHIR with CCDA and security aspects (leveraging OAuth2 and OpenID Connect).

## Value Proposition

From the points above, the main value proposition for FHIR becomes evident:

- **Cost-effective to implement**  
While it is unlikely that healthcare interoperability will ever be easy or low in cost, FHIR removes many of the common cost barriers.
- **Rapid to implement**  
With freely available libraries to assist in development, use of familiar tooling and the rapid adoption that we are seeing by vendors across the sector, moving data from one system to another will not require a long period of understanding – and then implementing – systems that talk to proprietary interfaces.
- **Builds on the experience gained in Version 2 and Version 3/CDA with relatively straightforward migration possible**
- **Flexibility**  
Solutions can be rapidly adopted, tested, and then changed as requirements also mature and change.
- **No vendor “lock-in,” which occurs when proprietary interfaces are needed**  
This will provide increased choice to consumers, and encourage vendors to innovate on features and functionality.

- **Incremental adoption**  
Because there has been considerable work to ensure alignment between FHIR and existing standards, it is not necessary to “rip and replace” to implement FHIR. Rather, implementing FHIR can be done gradually, allowing an organisation to gain experience as they go. Vendors offering support products such as integration engines that can convert between different standards will assist this process.
- **Support for emerging technologies such as mobile and devices – including the “Internet of things”**  
Common standards for representing and moving data around can only assist their development.

While a similar list has been made by other standards, the key difference is that by involving implementers from the beginning of the work, issues that could impede uptake are recognised and corrected early in the process; thus, further enhancing uptake, and helping implementers see that their needs are being taken seriously by the standards community.

## The Ecosystem

In the context of healthcare information, an ecosystem refers to all the different components involved in healthcare, and more specifically, the ability to share information quickly and safely between them. FHIR certainly facilitates that by defining the patterns of exchange, and to some extent, the model of information (at least its representation “on the wire”).

FHIR is not sufficient by itself; it needs other components – particularly in the security domain, such as OAuth2 and OpenID Connect – which is why projects like SMART and Argonaut are particularly important.

The diagram below shows one way that consumers can access data (read and write) from many different systems by using these common standards.

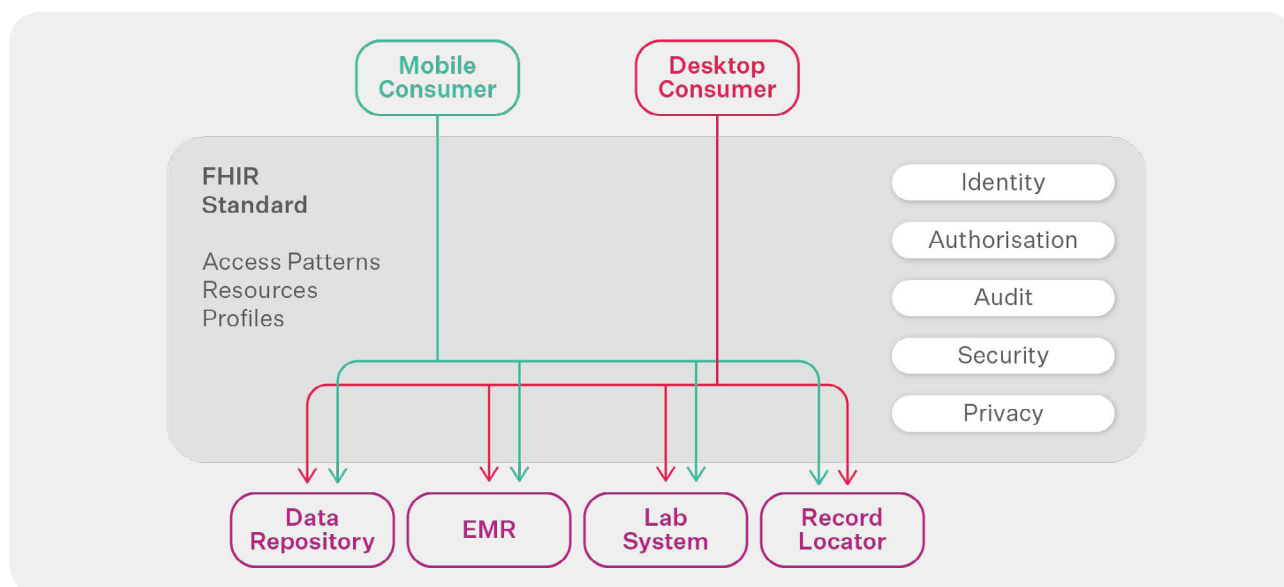


FIGURE 4: FHIR Access Example

## What FHIR is not

Despite all its advantages, FHIR is not a panacea for interoperability. There are many other requirements that need to be met, particularly in the way that information content is modeled in detail, security/privacy requirements, terminology and the workflows/behaviors that are required for quality healthcare. FHIR makes some progress towards these, and offers “hooks” that other standards can use, but there is still development needed in the way that FHIR works with those standards.

Grahame Grieve (the original author of FHIR) has a saying: “You can move complexity around, but you cannot completely eliminate it.” FHIR resolves a number of the common issues that implementers traditionally grapple with, leaving them free to focus on the more complex issues of representing clinical practice and behaviors, such as designing care plans and engaging in other aspects of shared care.

## Looking forward

At the time of writing this article, FHIR is in a Draft Standard for Trial Use (DSTU) state, which means that while it is in a fit state to be used, it is not yet complete and needs real-world use to ensure that it becomes comprehensive enough to cover all interoperability requirements. The upside of participating in using FHIR at this time is that doing so provides the opportunity to influence the standards development and to ensure that any potential individual requirements are met. The downside is that the specification will change, and that implementers need to be ready to accommodate those changes.

We expect that FHIR will be in a “normative” state (by which time, future releases will be “backwards compatible”) around 2016/17.

What is the future of FHIR? It is likely that FHIR will become like the Internet: ubiquitous and largely invisible – it’s just there. No one questions the use of the Internet and the standards that support it; it seems probable that that could be FHIR’s future also – as the “lingua franca” of healthcare data exchange.

## Why start now?

The previous section might make readers ask, “Why start now? Why not wait until FHIR is a normative standard?” The reason it is a good idea to get involved now is:

- FHIR solves real issues that people have right now, and for which there is no modern standard – especially in the mobile and devices space. Yes, an organisation can “roll their own,” but by using FHIR, you can take advantage of all the thinking that others have done – and this is considerable.
- By getting involved with the community, you can be sure that the specification will evolve to meet your needs, both now and in the future – and, incidentally, make the standard even better by your input.
- It takes time to “re-tool” to implement a standard – even one as straightforward as FHIR. Getting involved now gives an organisation the time to make necessary modifications to existing systems, and means your customers are likely to see you as forward thinking and innovative.
- How to get started

## How to get started

So what is the best way to learn more about FHIR, and to actually get started?

The specification itself is freely available on the web and fully hyperlinked (as we have come to expect with any web-based document), making it surprisingly readable, even to those who are not used to healthcare standards. The link is here: <http://www.hl7.org/implement/standards/fhir/>. There is also a Wikipedia page that has a great deal of information including a number of blogs dedicated to FHIR: <http://wiki.hl7.org/index.php?title=FHIR>.

A practical way to get started, especially where there is an existing HL7 v2 infrastructure, might be to use an Integration Engine such as Rhapsody™ Integration Engine, to map between HL7 v2 and FHIR, as shown in the image below (and described in this blog post: <http://fhirblog.com/2014/10/06/fhir-messages-part-2/>

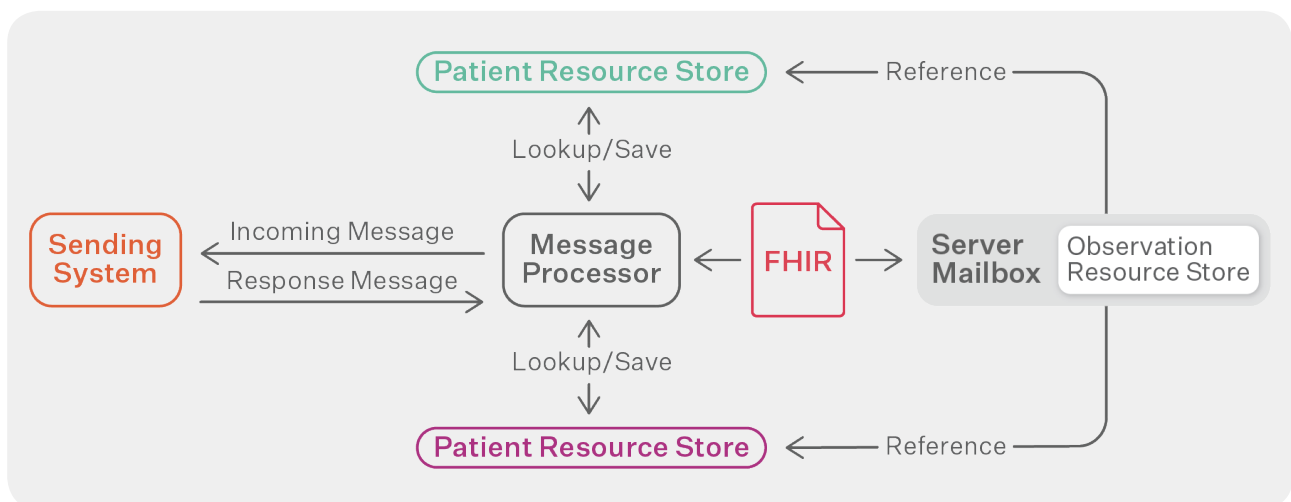


FIGURE 5: Map Between HL7 v2 and FHIR



## Conclusion

At the beginning of this paper, we described some of the challenges facing healthcare, which can be summarised as an increasing demand on the system by a population that is increasing in size, increasing in age, and developing conditions that require long-term (and costly) management. The workforce is not growing at a rate that can manage this; therefore, involving patients in their own care is becoming more and more critical. FHIR can help with managing these challenges by:

- Removing the common barriers to interoperability, making information available to those who need it and therefore, maximising the value of time spent with the patient and reducing waste from duplication.
- Allowing implementers to focus on the more complex issues of clinical behavior, rather than the “nuts and bolts” of representing and moving information around.
- Exposing information with widely used open standards, which enables the “clinical ecosystem,” where innovators can interact with data that was previously unavailable in proprietary silos of information. This will make it much simpler to develop applications focused on particular areas, especially in the mobile area.
- Standardising how that data is represented (in conjunction with terminologies such as LOINC or SNOMED), making it straightforward to apply more advanced technologies such as Clinical Decision Support.

## 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 Organisation (HISO) committee in New Zealand, which provides technical advice on standards to the National Health IT Board, and frequently writes about FHIR on his blog at <http://fhirblog.com>

