Enabling the Ecosystem with FHIR

Orion Health White Paper
Dr. David Hay
Product Strategist
112017

ORION HEALTH

How can healthcare information be made available, where and when it is needed?

The modern era has provided easily accessible healthcare information through the collection and manipulation of data, but it has also created complexity. We see this in the amount and type of data that is available, the growing number of sources where it is captured and stored, and the more specialised ways in which it is being used. There is the emergence of 'personalised medicine', where advanced analytics can be applied to this information – including the person's genome – giving management advice that is tailored to the individual rather than what has previously worked within a similar population. Following that advice often requires access to highly specialised services, but finding them can be a challenge.

And this material is complicated – it requires a deep understanding of specialised domains, vast quantities of information about both an individual and a population and the development of complex algorithms. With machine learning, facts are gleaned from this morass of data requiring powerful computing, more than a 'normal' Electronic Health Records (EHR) application can provide. How can this information and these capabilities be made available when, where and to whom they are needed?

HL7® FHIR® or Fast Healthcare Interoperability Resources can enable an 'ecosystem' supporting access to, and manipulation of, healthcare data and services by many different applications. In this White Paper we'll take a closer look at how FHIR and several derived standards could underpin such an ecosystem, and the types of services or components that will be needed.

FHIR is a healthcare interoperability standard which covers two main areas:

- A content representation in the form of 'resources' that describes how to assemble information into 'packets' that can be exchanged between IT systems.
- Definitions of how to actually exchange those resources in a variety of ways, as documents, messages or via a real-time Applications Programming Interface (API).

We also need a common understanding of the things or concepts we want to share in addition to the actual mechanism of sharing – if FHIR is the letter and the mailbox to use to send it, then a Terminology is the language that we use in the letter. Systematised Nomenclature of Medicine (SNOMED) and Logical Observation Identifiers Names and Codes (LOINC) are commonly used terminologies in FHIR deployments.

How do these concepts merge, what do we mean by an ecosystem, and why it is a good thing?

Wikipedia describes a digital ecosystem as: "A distributed, adaptive, open socio-technical system with properties of self-organisation, scalability and sustainability inspired from natural ecosystems".

Some of the key words in this definition as they apply in our own context:

- Distributed – information is stored in multiple different places, rather than a single large repository.
- Adaptive – it is able to change to meet new requirements.
- Open – we'll take this to mean the use of standards that lower the barrier to entry. For example, there isn't a cost to actually use the standards, and as the participants in the ecosystem implement these standards it is straightforward (rarely easy) to talk to each other.

- Self-organisation – once the components (at least their functions and interface) are defined, there isn't the need for some over-reaching coordinating system.
- Scalability – as more participants join the ecosystem, performance is not adversely affected (within reason).
- Sustainability – the ecosystem can 'run itself' without the need for on-going support (again, within reason).

So how big is an ecosystem? And is there more than one in the world?

For our purposes, we're going to assume that the ecosystem has some pre-defined size (in terms of the coverage area) – for example a city or a country – or maybe a particular specialty, as it's unrealistic to assume that it will be global (initially at least). This is not so much a consequence of the technology, as that different countries have different requirements and different operating processes.

At a minimum, we'd expect that even participants from outside the ecosystem can access the data within it (subject to security and privacy constraints of course) even if they are unable to fully participate.

And why is it a good thing? Well, it's all about the ability to easily share and exchange data between participants. To maintain and improve the health of individuals and populations, access to information and services is paramount. Whether concerning the delivery of care to individuals, decision support or analytics against populations, data is a fundamental requirement, so we want as low a barrier to participation as possible – whether the barriers are financial, technical or political. Ecosystem thinking gives us the best shot at that.

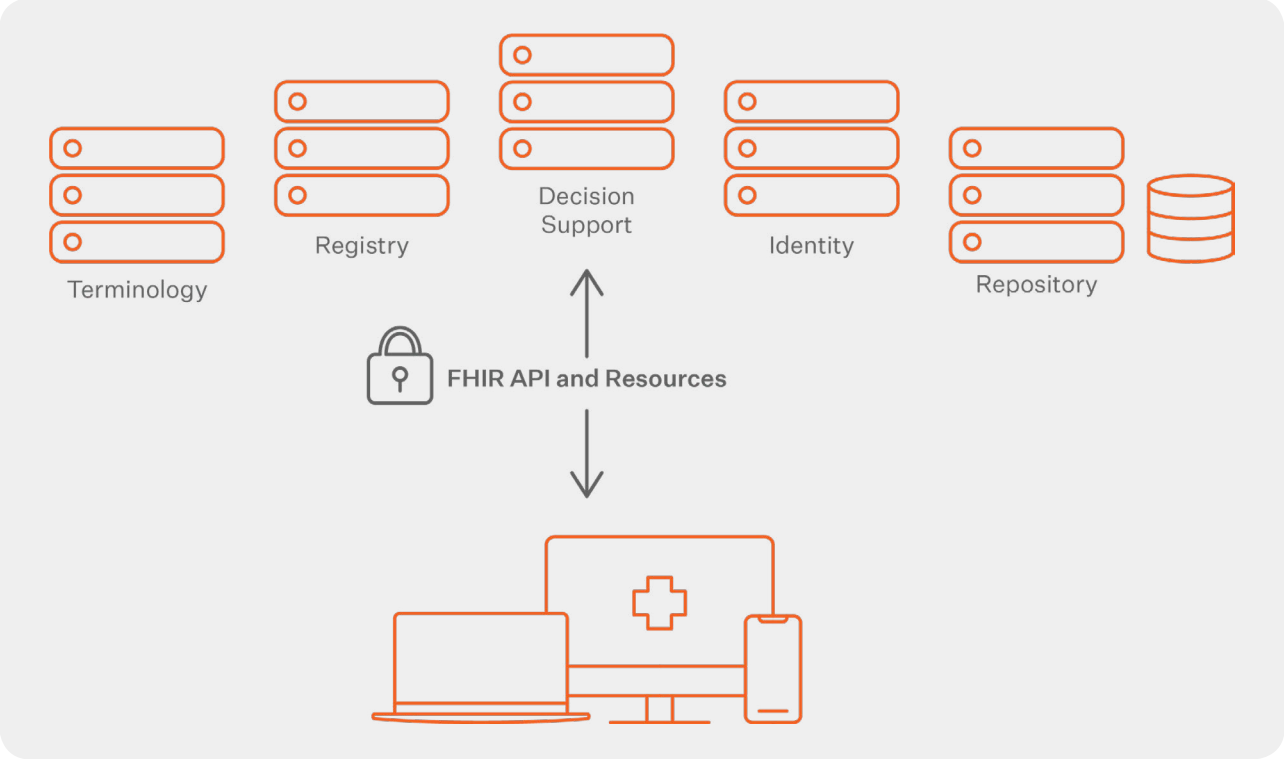

FIGURE 1: FHIR API and Resources

Here's the 'roadmap' for our discussion:

- Paradigms of exchange
- Defining the content to be shared – logical and physical
- Terminology servers
- Identifying the participants
- Patient data / Repositories
- Workflow services
- Decision Support services
- Security and Privacy
- Management and Governance

Paradigms of exchange

There are a number of different ways that we need to share information. Broadly, we can divide them into 3 main categories – Messages, Documents and Real-time.

Messages are used when one system wishes to inform another of an event, or is requesting that they do something. A key characteristic of a message is that once it has been delivered, the message itself does not need to be retained (other than for support and audit purposes). An example where messaging is appropriate is when a person is admitted to a hospital, and the hospital wishes to inform others of that event.

Documents are appropriate when summarising information for a purpose (generally clinical), and when that summary needs to be retained indefinitely. Examples of documents are a Discharge Summary or a Progress Note. Also, documents do not generally directly cause actions to occur.

Unlike messaging or documents, **Real-time** interactions occur across an **API** when a person or application is directly interacting with a system and expecting an immediate response from the system. An example could be a patient checking what medications they should be taking, or a clinician updating consultation notes and creating a clinical plan. Real-time interactions are commonly referred to as **RESTful** interactions (which describes the actual architecture), though other architectures are possible.

Our ecosystem needs to support all of these paradigms, which implies a number of supporting services we identify in this paper.

Defining the content

Regardless of the exchange paradigm, we need to think about what is the information being exchanged, and what it looks like 'on the wire'. There are a few aspects to this.

First we need a language that both sides understand - this is the Terminology that we referred to at the beginning. FHIR is not itself a terminology, rather it refers to one when it needs to describe a specific concept like a procedure or a diagnosis. Common terminologies are SNOMED and LOINC, but there are no inherent restrictions, you can even make your own if needed (though this isn't a recommended approach).

Next the representation of the actual information items within the exchanges needs to be defined. For example, consider a prescription - a request for a medication to be given to a patient. As well as identifying the patient and the prescriber, we need to identify the medication, why we are prescribing it, how much and how often it should be taken, for how long, with or without food and so forth.

We need to define both the information in the interaction (sometimes called an Information or a Logical Model), and the representation 'on the wire'. The former can be used by the Clinician to define 'what' it is they wish to share, the latter is the 'how' it gets communicated.

But different jurisdictions (e.g. countries) have different content and requirements for this. There might be a different terminology used to identify the medication. Some medications may require a specific authorisation or a reason to be given or an indication whether the dispenser can supply a different (but equivalent) medication.

In FHIR, we use 'profiling' to describe these differences from the base concepts We take the base resource (in this case Medication Request) and then remove the parts we don't want, add the additional ones, and set any terminologies that are specific to our use.

So our ecosystem will need a profile registry to hold all these profiles so they can be accessed as needed. (We refer to this as a 'Conformance' registry as the artifacts involved are called 'Conformance' artifacts.

Given that this profiling work is an ongoing process, it would be good if the ecosystem provided tooling to support it. Clinical Design services make it easier to co-operatively design a new profile, get feedback on them, publish and maintain them. We may choose to use the more clinician friendly information models as the basis for the discussion, and generate the profiles directly from them (either automatically or by a FHIR expert). This will make it easier for Clinicians and Business Analysts to understand and to participate in the design process.

Finally, there is a workflow process that we'll need to manage for electronic prescribing (and others) which will be discussed later in the paper.

Terminology Services

A lot of the information is going to be coded in some way. Coded data is better than simple textual data both to ensure shared understanding, but also when itc omes to providing Decision Support at the point of care (increasingly important with the huge amount of data and increasing understanding of that data that we are experiencing), and also when it comes to the secondary uses of the data – population based analytics for example, or as input into personalised care systems.

So we need a way to make it easier for the providers of information (Clinician and Patient) to code data in a consistent fashion. FHIR provides a number of components that can help here, ValueSets to list the preferred terms to use in a given scenario, and a number of Terminology operations to make it easier for applications to provide User Interfaces for the users (such as pick-lists). So let's include a Terminology server in our ecosystem that everyone can use.

The following diagram shows an application using the Terminology Service to retrieve data to display to the user. The complexity of the operation (and it can get exceedingly complex) is performed by the Terminology server, and provide to the consuming application via a simple interface.

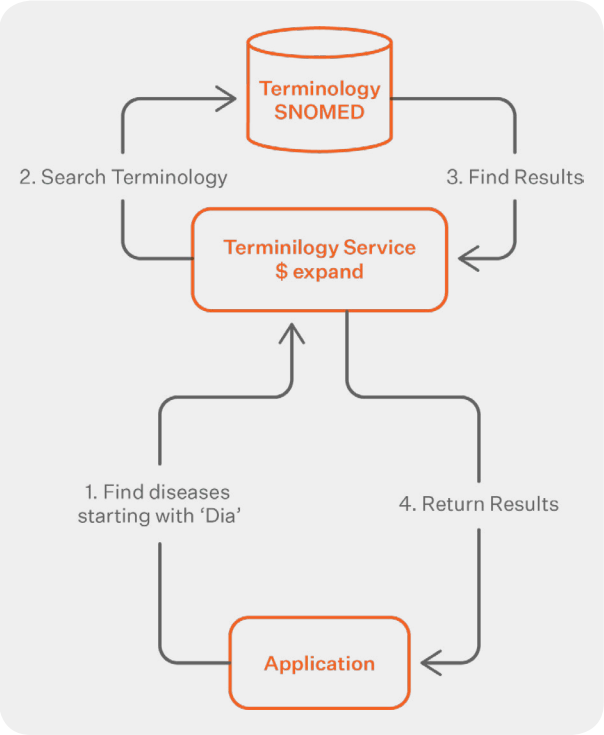

FIGURE 2: Example Terminology Service

Identifying the participants and locating services

By definition, our ecosystem is going to have a large number of different participants – human, device and application. We need to be able to identify them for a number of reasons:

- Being able to consistently refer to them is critical. In order to be able to locate all the information about a particular patient, every system storing data needs to use the same patient identifier (or, at least, there is a way to say that two identifiers refer to the same individual).
- We need to know who has authored a specific piece of content, who can access it – and, indeed, who has accessed it.
- We also want to identify the 'types' or categories of services that are available and have been delivered – important for managing service delivery.

So we're going to need a number of registries:

- A **Patient registry** that identifies the targets of healthcare delivery.
- A **Provider registry** that identifies the people who can provide care. This will need to include things like what skill set a provider has, where they are available, specific permissions they may have and so forth.
- A **Services registry** that details the services that are available, and where that service can be accessed.

These registries are going to be especially important for some of the workflow requirements we have – referring to services, ordering tests.

But there are other registries that we may want. For example:

- An Application registry that identifies specific applications (like a mobile app for managing diabetes) that we have checked via a certification service. We do want to be sure that applications displaying or updating the shared data are doing so correctly – and safely.
- A Devices registry – e.g. glucose meters that have been tested and shown to be accurate.

We're also going to need a way to find these registries. So either they are 'well known' – we know where to find them – or there is a 'meta-registry' a registry of registries!

Getting Patient specific data Of course, the ecosystem exists to share clinical information, so being able to find and use the clinical repositories where that data is stored will be needed.

As a baseline of course, the use of FHIR largely determines the structure of the API that the repositories expose, but individual servers can choose what part of that API they support – and can also define additional interfaces – such as custom operations and search parameters. Fortunately, FHIR defines how a server can make this information available (the CapabilityStatement resource) so if you know where the server is then you can determine its capabilities, but there are two issues to resolve – where are those servers and which patients do they have information on.

To know where the repositories are, we need a **Repository Registry** (though we could extend the **Application Registry** for this) finding all the data for a single patient is a little more complex.There are a number of strategies we could follow:

- If there aren't too many of them, then it might be feasible to just call all of them when looking for a patient's data. We could put an aggregation layer in front to make it easier for the client.

- An alternative is a **Record Locator Service** which holds an 'index' of patient and repository (similar in concept to the IHE XDS profile). This could range from a simple 'this repository has data about that patient' to a more complete index of all data in the ecosystem.

We also need to think about the nature of the clinical data we are recording. In some cases, it doesn't matter if more than one repository has data about the patient (leaving aside the issues of duplication) – by way of example, if there are laboratory results or encounter information in different repositories we will have to obtain them.

But it's a different matter for 'curated lists' such as a medication list or a list of allergies. Particularly with medications, if there are lists for the same person on different repositories then there are very real clinical risk issues. Which medication list is the correct one to look at? If a patient wants to look at their definitive list, where do they go? In these cases, we may want to have just a single repository for medication lists, or if there are multiple repositories then a single patient uses only one of them.

Care Plans are between these two extremes. FHIR allows there to be multiple Care Plans for a patient, with the overall plan being the 'union' of them. But understanding all the planned interactions between a patient and the ecosystem is a fundamental part of efficient healthcare delivery, so we may choose to have a 'master' copy of the plans in one repository per patient, or allow aggregation to occur.

Workflow services

Workflow can be considered as a co-ordination requirement – tracking actions that need to be performed both to ensure that they are completed, but also to make them visible to others.

For example, consider referrals. We could make available in our ecosystem a **Workflow Service** that is used to manage all referrals. When a referral is made, the details of that referral are stored in the service, and as thereferral proceeds it is updated with the current state. Then whenever the patient is seen, care deliverers are aware there is some investigation under way, and it is easy for the participants in the workflow to update the service.

The workflow service could be used for any process that has multiple steps, such as appointments, the prescribing and dispensing of medications, or the ordering of laboratory tests and other investigations (though we might want to split out the appointments into an **Appointment Register** and also have a **Contact/Encounter Register** for completed appointments). This could support more innovative processes, for example we could use the provider registry to locate a suitable target for referrals based on the recommendations from a Decision Support service, and an Appointment service to perform the scheduling taking into account existing appointments that the patient might already have.

Decision Support Services

As mentioned at the start of this whitepaper, medicine is complex, and becoming even more complex as our understanding improves and treatment options expand. Already it is impossible for any single individual to be aware of all the possibilities, so Decision Support Services are going to be a must for our ecosystem. We can think of these in two categories, those that can operate without accessing the patient record, and those that do need patient specific data.

As an example of the first, consider an Immunisation Recommendation service.

Especially for children – but also for the older person or someone with a chronic condition – there are a set of immunisations that are recommended at particular ages/stages. These recommendations can vary, and it can be tricky to decide exactly what an individual should receive especially when changes occur and a 'catch up' program is required. So imagine a service that you call with a person's Date of Birth and gender plus their allergies and immunisations they have already received and get back a list of what their immunisation program should be. Or, what vaccinations someone visiting another country should receive. Simple to build and not containing any Personal Health Information (PHI).

An example of a service that does require existing patient data would be the 'Personalised Medicine' applications being built around the world. These applications use complex algorithms and machine learning capabilities to examine a record and make management suggestions. These applications can expose their services using the Clinical Reasoning Interfaces being defined in FHIR.

One new standard that we will likely want to support is CDS-hooks. This is a standard based on Substitutable Medical Apps and Reusable Technology (SMART)® and FHIR that allows an application to automatically invoke a service, passing across requested data, or allowing the service to access data within the application and use that as the basis of recommendation. For example, when writing a prescription, the application automatically invokes a Drug Interaction service that passes across the medication being prescribed and allowing the service to access the person's medication and problem history from within the EHR database (or from some other source of clinical data). Because the interface is standardised, there could be a number of different services that a user could choose from, or that the ecosystem might make available.

Here's an example of CDS-hooks in operation. The diagram is based on the CDS-hooks specification:

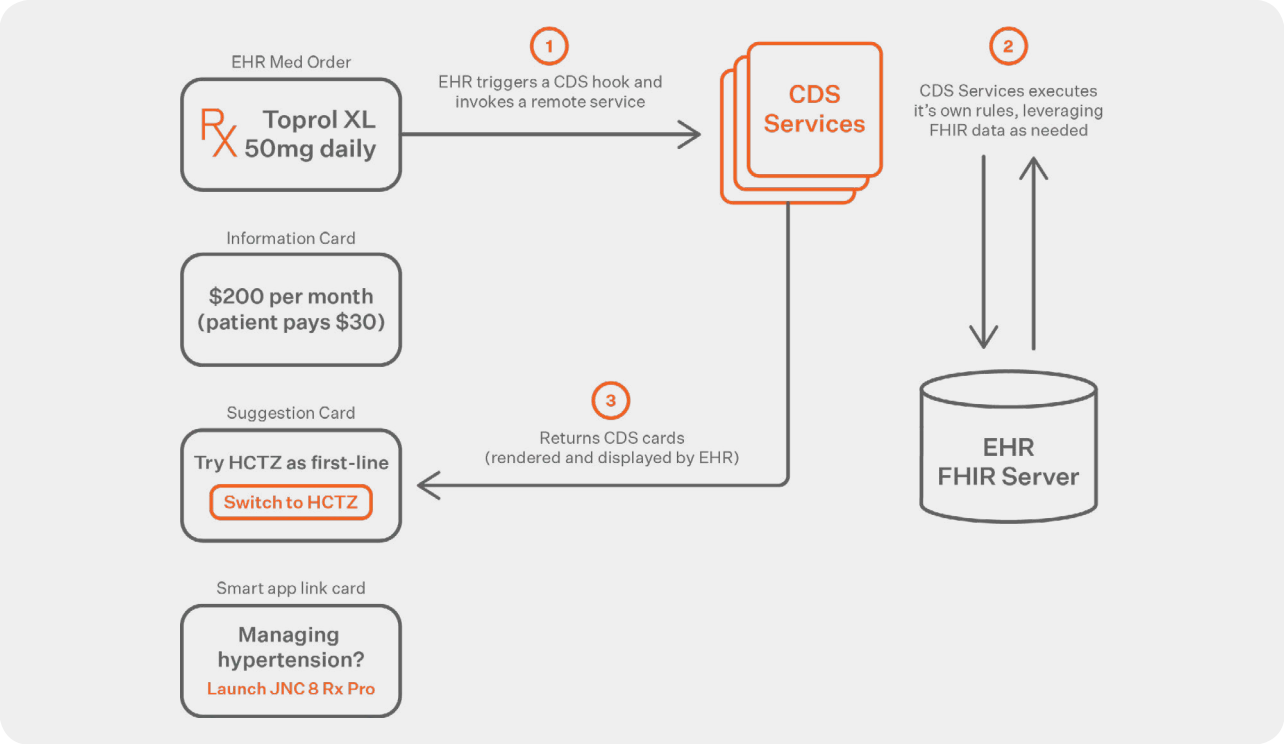

FIGURE 3: Example of CDS-hooks in Operation

Security and Privacy

Given the sensitive nature of the data we record in healthcare, robust security and privacy mechanisms will be paramount. It is well beyond the scope of this paper to describe them in any detail and the requirements vary tremendously across jurisdictions, but the things we'll need to cover off include:

- Use of encrypted communications is fundamental – whether Transport Level Security (TLS) for HTTP traffic, or encrypted messaging/document sharing with certificates.
- Access controls – Oauth2 and **SMART.**
- Security policies – especially around the handling of sensitive data.
- Robust auditing. We may choose to set up a common **Audit Service** which logs all record accesses within the ecosystem (based on IHE ATNA for example). One of the requirements we have is that the patient needs to know who has accessed their record, and this could be the basis of providing that functionality.
- Managing consent information, including access to data for care delivery, research, Advance Care directives and so forth.

Incidentally, this is an area where blockchain technology may play a part. Best known for its use in the Bitcoin financial system, blockchain essentially establishes a way to record transactions in a way that is non-repudiable, so that it cannot be denied. The exact place for blockchain in healthcare is unclear at the moment, but the recording of consent and audit access would seem to be a good fit at the least.

Management and Governance

Although the ecosystem is described above as 'self organising', there is still going to be a layer of governance that is needed to define and choose the standards, and ensure that they are followed by the participants. We do want to make it as easy as possible for participants to do 'the right thing', so there are some things we need to do. We need to manage who has access to the data in our ecosystem, this is highly sensitive data, so strong controls are needed.

It needs to be easy for potential participants to find the information they need. The **Conformance Registries** will help with this, and of course the FHIR specification is an on-line resource, but we will also need to create and expose **Implementation Guides** that have the details of our specific implementations including examples and other pertinent documentation such as Use Cases and Scenarios.

We should set up development sandboxes with sample data that developers can use to build applications, and Certification services than can validate that their solutions are correct. In an ideal world, these will be as automated as possible, though it does seem likely that some human interaction will be needed. Of course, we have our **Applications registry** to record apps that are compliant to our requirements. The whole process of profiling needs to be open and available for all to comment on.

Conclusion

So to summarise, our ecosystem is going to need a number of different registries and services, including:

- An agreed definition (model) of the information we're going to share expressed as FHIR profiles with the supporting conformance artifacts (resources) openly available
- A FHIR conformance registry that holds those artifacts
- Clinical Design services
- A Patient registry that holds the patient identifier
- Record Locator services
- A Provider and Services registry
- Applications registry
- Terminology Services
- Recommendation and Decision Support Services
- Scheduling services
- Workflow services
- Audit Services

And we're going to rely on the following standards:

- FHIR
- CDS-Hooks
- Oauth2 and SMART
- Security standards
- Terminologies

The goal is to enable the ecosystem with FHIR, to ensure healthcare information be made available, where and when it is needed. So that all participants even those from outside the ecosystem can access the data within it and they are able to fully participate within the ecosystem. It is important that all participants can easily share and exchange data between them. This will improve the health of individuals and populations, by providing access to specialised healthcare information and services. Whether concerning the delivery of care to individuals, decision support or analytics against populations, data is a fundamental requirement, so we want as low a barrier to participation as possible – whether the barriers are financial, technical or political. Ecosystem thinking gives us the ability to achieve a truly open healthcare data system.

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 patient's health information globally.