HL7 FHIR for Executives

Summary

The FHIR (Fast Healthcare Interoperable Resources, pronounced “fire”) standard is created by HL7. It supports the exchange of data between software applications in healthcare, combining the best features of HL7’s existing v2, v3 and CDA product families while leveraging the latest web standards and applying a tight focus on implementability.

In summary, FHIR:

- Is easier and cheaper than other comparable standards:
  - It is faster to learn, implement and troubleshoot (you should be able to “figure it out” over a weekend and the standard is shipped with toolkits and examples).
  - It has a vibrant and open source community and has frequently held connectathons.
  - It uses modern technologies, the same as used by e.g. Facebook, Twitter and Google.
  - There are more people familiar with these technologies (thus less expensive consultants).
- Is being implemented right now:
  - US: ONC, SMART, Intermountain, CommonWell. UK: NHS. NZ: Orion Health. NO: Helse Vest.
  - IHE Profiles
  - 70 implementations, 20 countries (2014)
- Is likely to significantly impact Health IT:
  - scales well from simple to complex
  - flexible
  - free and fully open

Application of FHIR

FHIR is suitable for application in a variety of settings:

- the classic in-institution exchange of data between systems
- in a regional setting (Regional Health Information Organizations)
- on a national scale, e.g. in national health hub’s or EHRs
- in social media and mobile applications

Why FHIR?

There has been a need to share healthcare information electronically since the introduction of mainframe based EMRs in the 1960s. HL7 version 2, aimed at the exchange of data between fixed departmental systems within the context of a hospital organization, has its technical roots in the 1970s. There is increasing pressure to broaden the scope of sharing across organizations and disciplines, mobile and cloud-based applications and to achieve integration in days or weeks instead of months or years.

These are the drivers of FHIR:

- **Shift in healthcare**
  The patient is in control and owns his medical data. There is an increasing pressure to broaden the scope of sharing health data and a growing need for exchange of information across organizations, disciplines and (regional) borders.
- **Shift from off-line to on-line**
  Recent years have shown the shift from PC to tablet, from web to app, from electronic health record to personal health record, from desktop to cloud. FHIR is lightweight, is made for mobile and it enables the data to travel with the patient.
- **Shift towards data transparency**
  EHRs and other medical systems tend to behave like black boxes. Data sits in there, you can't get it out - and if you can, the data is useless because it's incompatible. FHIR acts as an 'Open API' to access data in these silo-like EHRs. Different aspects of patient data will end up being hosted in different systems. New healthcare tools should have the ability to reach out to these systems and use that data in a collaborative manner.
- **Shift towards analytics**
  Analytics requires data transparency, but also for the data itself to be in a format which is optimized for analysis. FHIR uses data structures that allow one to easily slice and dice the data for analytics. Unlike CDA, with FHIR there is no need to split documents into more atomic concepts for analytics.

Start from scratch

The initial FHIR developers posed this question: What would healthcare information exchange look like if we started from scratch?

A web search for success markers of modern implementation approaches led to REST-based APIs. FHIR is a healthcare exchange API based on this approach, which provides a simple and efficient way to discover and consume information across distributed systems.

HL7, other standards organizations and the United States President's Council of Advisors on Science and Technology (PCAST) came to the conclusion in 2011 that the base unit of exchange when it comes to data should neither be too large (yielding unwieldy and overly complex data structures) nor too small (it should convey meaningful data). FHIR uses small logically discrete units of exchange with a well-defined behavior, meaning and contextual metadata.

The design philosophy behind FHIR

FHIR doesn't solve the really hard interoperability challenges – such as workflow and policy differences between organizations, and variability in the data captured and found in different contexts of use. FHIR aims to make the implementation of the data exchange as simple as possible, to clear the way for solving these hard interoperability issues.

The design philosophy of FHIR can be summarized as follows:

- **Focus on Implementers** – easy to understand for developers, plenty of tools around, APIs and examples available out of the box.
- **Target support for common scenarios** – but support the concept of extensions.
- **Leverage cross-industry web technologies** - use the same cross-industry technologies as Google, Facebook and others (e.g. XML, JSON, HTTPS, OAuth).
- **Require human readability as base level of interoperability** – as a fall back option for applications unable to interpret structured content – a lesson learned from CDA.
- **Make content freely available** – FHIR has an 'open source' license.
- **Support multiple paradigms & architectures** – FHIR leverages the same data models and profiles (see details below) everywhere regardless of interoperability approach (REST, Documents, Messages, Services). These are lessons learned from HL7 version 3, where different models are being used depending on the integration approach, leading to additional implementation challenges.

Building blocks of FHIR

The main building blocks of FHIR are Resources, References and Profiles.

Resources

Resources are small logically discrete units of exchange with a defined behavior and meaning. They are the smallest unit of transaction. The 150 different resource types cover all of healthcare. Examples include Patient, Practitioner, Allergy Intolerance, Family History and Care Plan.

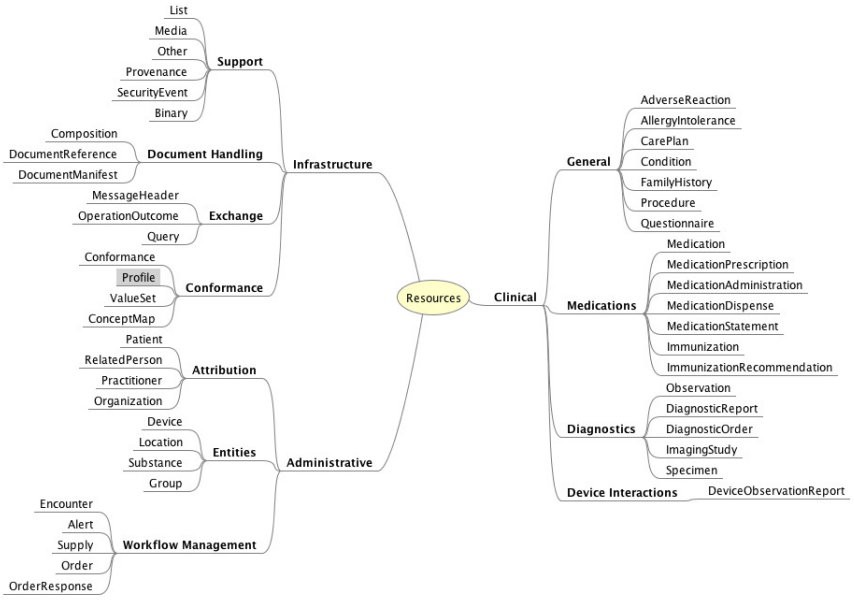

Mindmap of currently defined FHIR resources (2014)

A resource consists of three parts:

1. Structured data – attributes to support the 80% common use cases. "We only include data elements if we are confident that 80% of implementations maintaining that resource will make use of the element." Other content is pushed to extensions.
2. Narrative – textual summary of the content of the resource.
3. Extensions – attributes to support non-common use cases.

References

References are links from one resource to another. Using these, the resources combine to create a network (or web) of information that represent a health record (or at least, a useful part of it). Systems can navigate these links to decide what resources they need for a given task.

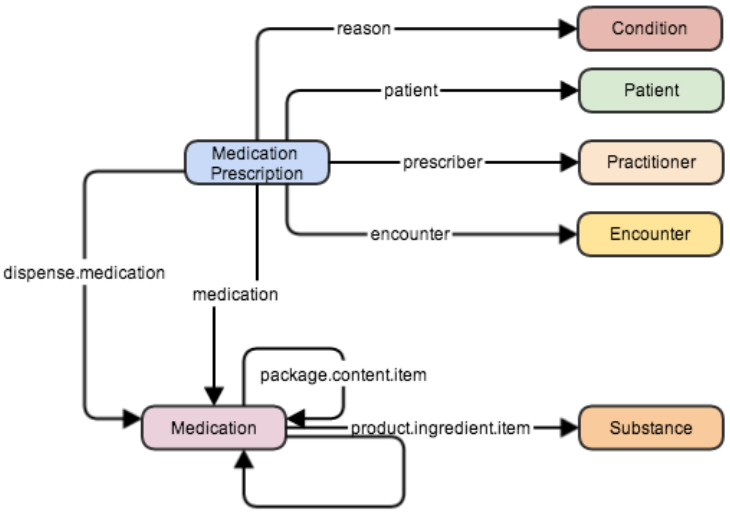

Example of references between FHIR resources

Profiles

Resources have no built in restrictions on how they are used - that job is handled by Profiles. Parties exchanging data define the specific way they want to use resources and their relations by using Profiles. Profiles are the framework for defining/discovering services. They define what a particular jurisdiction needs to communicate and store when it comes to Resources and their extensions. Profiles are defined by HL7, a country, a region, an organization or a project.

For example: A given hospital may have rules for the content of a referral for pediatric surgery at that hospital. It might require information such as primary care physician, parent names, child's age and gender, etc. If any of those pieces of information are missing, it wouldn't be considered a valid referral (for that type of referral at that hospital). A profile defines rules about the standard resources, saying things like: Which elements are required, what terminology codes are used by a particular system, and what additional extensions are defined.

We already use HL7v2 and CDA – why should I use FHIR ?

FHIR allows for the exchange of one and the same set of resources, irrespective of whether those are exchanged as individual resources, within documents, using services, or within messages. The United States President's Council of Advisors on Science and Technology (PCAST) report (2011) stated that "We think that a universal exchange language [for health data] must facilitate the exchange of metadata tagged elements at a more atomic and disaggregated level [than documents], so that their varied assembly into documents or reports can itself be a robust, entrepreneurial marketplace of applications."

FHIR adds a new architectural approach when compared to HL7v2 and CDA: a RESTful approach, which allows for individual resources to be retrieved, updated and queried for – necessary in the case of registries (e.g. patients, or pharmaceutical products) and which has the benefit that one can choose whether or not to retrieve a referenced resource. In HL7v2 and CDA one is forced to retrieve a collection of data, irrespective of whether one requires only partial data.

Example Resource, the unit of exchange in a RESTful environment

HL7 version 2, based on 1970s legacy EDI messaging standards, works relatively well within institutions. It doesn't scale well across organization boundaries however. FHIR messages cover the same functionality as HL7 version 2. A FHIR message consists of a list of resources (comparable to HL7 version 2 message segments).

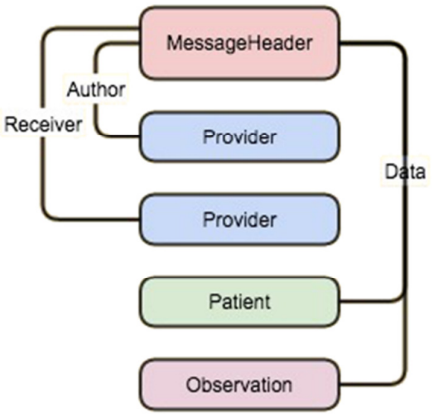

Example of a FHIR message with resources and references

CDA is broadly implemented, but has a very steep learning curve. The CDA standard has a dual purpose: to achieve human-to-human interoperability (text is mandatory) as well system-to-system interoperability (optional, software processable data). Interoperability beyond the human-to-human level is still a challenge however.

FHIR documents cover the same functionality as CDA: resources (comparable to sections in CDA) contain both text and software processable data. A FHIR document is a composition of section-level resources. The FHIR developers are working on a project (2014) to ensure that the contents of all sections of the US CCDA specification are covered by FHIR resources and their relationships.

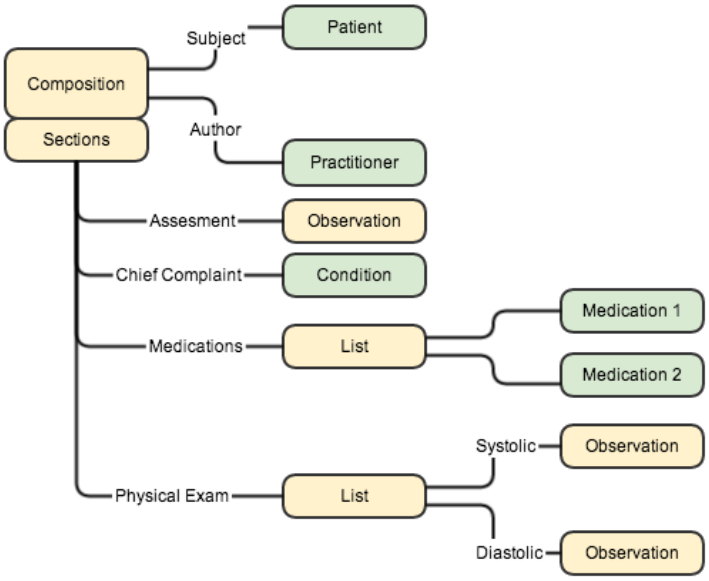

Example FHIR document with resources and references

FHIR currently fills a new niche using REST to support Social Web and Mobile Applications, and also has the potential to replace HL7 version 2 and CDA.

From v2 and CDA to FHIR

Most healthcare provider organizations already have to deal with multiple standards (e.g. HL7 v2, CDA, X12, DICOM) and mappings between them. The issue of having to map between FHIR and these other standards is no different from any the current mappings. Given that FHIR is based on HL7 v2 and CDA, and that there is a conscious effort to align the resource definitions with CDA, the mapping between the various HL7 standards will be relatively straightforward. Any of the existing interface engines should be able to support such mappings. Tooling for mapping of current standards to and from FHIR are expected to become available as the use of FHIR spreads.

Additional reading

For additional materials related to FHIR, we kindly refer you to:

- The FHIR standard (document) - http://www.hl7.org/fhir/
- Introduction to FHIR for Clinical Users (document) - http://bit.ly/1sYWE11
- FHIR for Executives (video) - https://vimeo.com/112905640
- FHIR elevator pitch for Software Developers (video) - https://vimeo.com/70111319

Summary

In summary, FHIR:

- Is easier and cheaper than other comparable standards.
- Is being implemented right now.
- Is likely to significantly impact Health IT.

Acknowledgements

Editor: René Spronk (Ringholm), based on materials created by Lloyd McKenzie (LM&A Consulting), Ewout Kramer (Furore) and Grahame Grieve (Health Intersections). We'd like to thank the various reviewers and contributors, notably David Hay (Orion), Rik Smithies (NProgram) and Rien Wertheim (Furore). Please sent any comments and/or suggestions to FHIR@ringholm.com.

Appendix: Who is already using FHIR?

FHIR is still a draft standard, expected to be normative late 2016. Nevertheless, we already see a huge interest in FHIR taking place. Vendors are spending R&D money on FHIR and sending their people over to conferences and connectathons.

See http://bit.ly/1svN172 for a list of organizations that are known to be implementing FHIR. Some example projects can be found below.

The CommonWell Health Alliance (U.S.)

The CommonWell Health Alliance is an independent, not-for-profit trade organization based in the U.S. The founding membership of the alliance are healthcare IT vendors (e.g., Allscripts, athenahealth, Cerner, CPSI, McKesson) . The goal of the alliance is to provide core infrastructure services for enabling healthcare data interoperability.

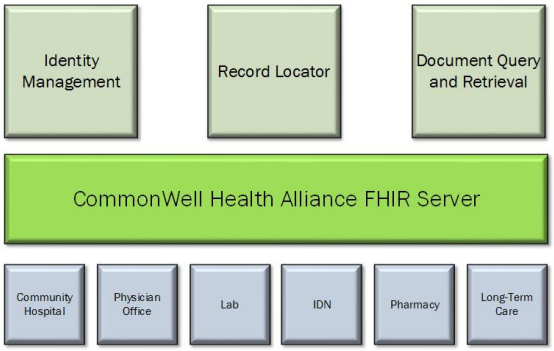

As the service provider for CommonWell, RelayHealth has implemented services focused on

1. Patient identity management: a REST-based patient link and matching API that supports patient matching by local patient identifiers and patient demographic data.
2. Record location and targeted query for clinical data: a REST-based service to locate patient records across the CommonWell network.

These services are based on open standards and, with respect to the REST services, utilize FHIR resource definitions. The FHIR standard was extended to provide a service-oriented mechanism for EHR systems to authoritatively "link" patient records across heterogeneous patient identity domains.

The pilot launched in early 2014 at hospitals and clinics in three U.S. geographies. The solution is now commercially available, and is expanding to new locations in the U.S. The development of the CommonWell service APIs is ongoing – e.g., adding FHIR-based document query and retrieve services.

For additional information, see: http://www.commonwellalliance.org/

Oridashi-Hiasobi Primary Care System (Australia)

In the Australian market 3rd party application vendors desire readily accessible interfaces and data representations for primary care system health record content.

The Oridashi-Hiasobi FHIR server provides a platform across multiple primary care systems in total comprising greater than 80% of the Australian market. The use of FHIR REST based read-only interfaces and resource representations creates a common methodology for application vendors seeking plug-and-play capability.

- Build once, many primary care systems.
- Remove cost of custom data interfacing.
- Focus on end user value, not data access.
- Open new markets for applications where FHIR servers are available.

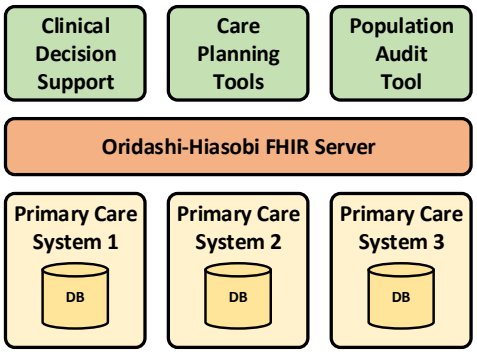

The benefits delivered with the FHIR specification, reference implementations, sample code and community engagement are compelling for application vendors integrating with FHIR servers. FHIR powered interoperability offers application vendors the ability to expand their market availability and drive down customization costs.

This product was first delivered to application developers in March 2014 as a core data source for green field implementations of a Clinical Decision Support and Care Planning Tool platform. These were delivered as a pilot in September 2014, and production use is currently in planning. Ongoing work on the same application platform will continue in 2015 to provide audit and quality tools for the clinic. These support aspects of the national Quality Improvement and Continuing Professional Development program managed by the Royal Australian College of General Practitioners as formal requirements of registration to practice in Australia.

For additional information, see: fhir.oridashi.com.au or email: brett.esler@oridashi.com.au

Mental health care data exchange (Netherlands)

Dutch mental health care providers are investing in e-health in order to support blended care plans, which mix face to face treatment with self-directed work in e-health interventions. These interventions are increasingly developed by independent companies and this provides an integration challenge for the existing EHR systems and e-health platforms.

Supported by health insurers, the Dutch mental health care sector has joined forces to develop a shared integration language and service to share data between e-health interventions and e-health platforms. This has the following benefits:

- patients can work more independently on improving their condition and self-reliance. It is also easier to involve their family and friends;
- mental health institutions have more flexibility in what they offer patients and will lower both the costs of labor and of IT;
- e-health platform developers can increase their market reach and broaden their product portfolios;
- health insurers obtain better care at lower costs.

FHIR emerged as the best standard to implement the architecture of what is now called “Koppeltaal” (which is Dutch for ‘Connectivity Language’).

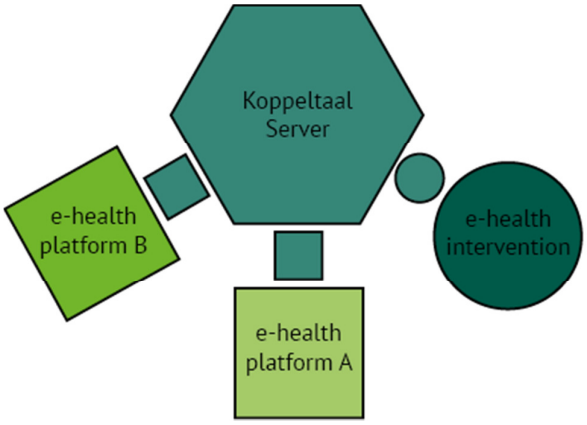

Applications can register with the Koppeltaal server and use a publish subscribe model to share data. In its first version, Koppeltaal supports the exchange of FHIR messages between e-health platforms and a game for children with Autism.

The key reasons for choosing FHIR are its built-in flexibility, its alignment with current internet standards, its extensibility, and its profile mechanism.

The first version is being tested as of December 2014 and Koppeltaal expects to be ‘live’ in the first quarter of 2015. After phase 1, Koppeltaal expects to extend the language and service to connect to EHR and Routine Outcome Measurement (ROM) systems.

For more information about the project contact Sergej van Middendorp at sergej.van.middendorp@milesahead.eu

Laboratory results in Helse Vest (Norway)

Helse Vest RHF is one of four regional health authorities in Norway and has the overall responsibility for the specialist health service provided in the western part of Norway. Helse Vest owns and manages the five health trusts in the region.

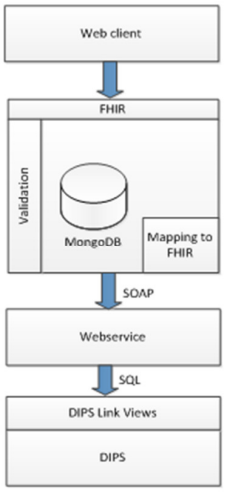

In 2014 a new requirement came up to share Laboratory results with several other systems (e.g. OR Planning, ICU and Cell Therapy). There are currently four different laboratory systems in use within Helse Vest, which will probably all be replaced by a new application in the next few years. Helse Vest doesn't wish to invest in the interfacing capabilities of these four legacy laboratory systems. The current EHR system (DIPS) orders and receives laboratory results and as such can be used as the basis for sharing information about the laboratory results.

Helse Vest chose FHIR as the most appropriate standard, mainly because of its open nature and its ease of implementation. The time spent on the creation of the interface was limited to about 2 weeks. The scope of the implementation is limited to the Observation resource and supporting resources such as Patient and Organization. The current Proof of Concept solution is expected to be put into production in early 2015.

For additional information, e-mail lars.gunnar@helse-vest-ikt.no