



INTEROPen

FHIR® and openEHR

Authors: Dr Ian McNicoll of openEHR Foundation; Dr Amir Mehrkar of INTEROPen; Dr Tony Shannon of Ripple Foundation

March 2019

Executive Summary

Poor interoperability between health and social care applications remains a critical barrier to the effective digitisation of health and social care. Whilst [HL7 FHIR[®]](#) is rapidly becoming established as the preferred data exchange standard in the UK, a number of NHS organisations and vendors are deploying applications built on the openEHR standard.

This paper outlines the differing approaches and goals of each standard, and suggests that both have a valid and complementary place in the challenges faced in digitising health and social care. FHIR[®] is more focused on data exchange between existing systems, whilst openEHR offers an alternative approach to building new systems and applications, which removes some of the barriers to interoperability in the first place.

An Ecosystem of Standards

Many reading this will already be familiar with the challenges of 'interoperability', the work of the [INTEROPen](#) community and the related [HL7 FHIR[®] Care-Connect](#) profiles.

Others will be familiar with the value of other organisations such as [SNOMED](#), [IHE-UK](#) and IHE standards such as IHE-XDS for document registration and location, and IHE-ATNA for audit logging.

Some will also have heard of [openEHR](#) but may be less familiar with the detail, and perhaps confused that it seems to be operating in a very similar space to FHIR[®]: clinical models, curation, data standards, and so on.

If we regard the overarching challenge as “how do we get this complex patchwork of applications to behave more like a single coherent system?”, then openEHR and FHIR[®] provide complementary but different, approaches. Primarily designed to solve somewhat different parts of the problem space, each offers a way of overcoming the challenge of capturing, accessing and exchanging healthcare information with minimal (*ideally no*) loss of meaning and context.

The seeming overlap of FHIR[®] and openEHR has elicited a number of questions from within the INTEROPen community, which we will attempt to address in this paper. These include:

- ***“Are NHS health providers and organisations expected to only use HL7 FHIR[®] for data exchange?”***
- ***“Why are some NHS health providers and organisation expecting to make use of openEHR as well as HL7 FHIR[®]?”***
- ***“Does openEHR ‘do’ interoperability?”***

In an attempt to provide fair and helpful answers to these questions, we start this piece with a summary of the health data exchange challenge (the bigger picture) and then turn our attention to how HL7 FHIR[®] and openEHR differ in purpose and use.

We contend that these technologies are best seen as complementary, not competing, approaches to solving some of the problems of health IT development, procurement and 'interoperability'.

IHE and SNOMED support other critical parts of the exchange ecosystem but their role is easier to define and both are complementary to HL7 FHIR[®] and/or openEHR.

The intended audience is CCIOs, CIOs, informaticians, and technology policy-makers. Just like at the [INTEROP SUMMIT](#), we hope to educate; this is important to us because part of our motivation to write this is to encourage a more open-minded and respectful dialogue amongst standards 'experts' in contrast to the tribal wars that unhelpfully play out on the Twittersphere and in niche groups. There is no need for this divisive behaviour; what patients and clinicians need is for us to all work closer, share our knowledge and experience and be brave enough to admit that these wicked health tech problems require us to establish a culture of learning, collaboration and trust if we are to make progress and innovate. And sometimes such progress may manifest in NHS Healthcare providers using openEHR and HL7 FHIR[®] in a blended approach, as well as SNOMED-CT and key IHE standards.

The bigger picture

It is not controversial to state that healthcare IT faces enormous challenges and that in spite of best endeavours by most of those involved, including commercial implementers, the rate of progress in healthcare digitisation is well behind that in other comparable sectors. In many respects this has the signs of a 'failed market', with the current UK GP systems market contraction being a localised example of a universal and global problem.

The focus of much of policy-maker concern is on achieving healthcare 'interoperability' i.e. being able to maximally exchange meaningful, computable, actionable data between different health IT systems. This is a pre-requisite for the digital transformation seen in other sectors, and in healthcare such interoperability can support both professionals and patients to deliver safe and effective care.

However, in spite of decades of investment, this has proved elusive, with different commercial, technical, regulatory barriers being variously blamed for the failure. Yet these barriers are fundamentally no different in other industry sectors, so what makes health different?

Arguably, what makes health different is the shape and scope of the data itself, reflecting both the complexity of human physiology, disease, and related healthcare processes and the unusually distributed, specialised and oftentimes siloed nature of much of health and social care practice. A clear example is the complexity of [medication management and dose timing instructions](#) when compared with financial ledgers.

In this complex environment, the purchasing community (national or local) tends to waver between two extremes, opting to either:

1. **acquire a single system**, generally from a single provider, which gives the advantage of much tighter data consistency and therefore less need for 'interoperability', but brings the downsides of potential vendor lock-in and often associated technical lock-in. As such "monolithic" systems grow and take on more and more care functions and services, their internal complexity can itself become a barrier to rapid innovation.
2. **adopt a 'best-of breed' approach**, acquiring multiple apps/systems from multiple vendors, then using 'interoperability' solutions to attempt to replicate the kind of tight data consistency offered by a single system; inevitably with significant complexity, cost and long development roadmaps.

In reality, at any one time, there tends to be a mix of both approaches, and even where a single system is intended to predominate, there are typically hundreds of small surrounding 'feral' apps, including quality improvement registries e.g. for cancer, cardiology, rheumatology or renal medicine, departmental and clinical hobbyist apps. These are rarely integrated with the main EPR system, but often grow to become a key part of service delivery.

In addition, 'systems' are nearly always purchased as whole entities: the application/UI, the associated internal clinical concepts, information models, workflow, guidelines, rules, and the database itself. All provided by a single vendor, all in lock-step and all completely siloed to that vendor (even if the code is open-source). This is the heart of the issue that 'interoperability' is trying to address: siloed data models/workflow and decision support rules. This is the 'interoperability problem'.

However 'interoperability' is not the whole problem facing our sector. Globally, there is an increasing dominance by a smaller number of suppliers, and in spite of significant investment in promoting 'apps development' and other innovation schemes, there has been a dearth of new entrants to the market able to grow to scale to compete with incumbent suppliers.

"There are over one hundred commercial suppliers of electronic health record software. Despite this, the market is currently dominated by a few large suppliers."¹

The current UK GP system market, and initiatives like [GP IT Futures](#) as a response, serve as good examples of this monopolisation problem; unfortunately it is also quite typical of other parts of the market globally.

There are many factors involved, but the challenges of understanding and storing health data, adapting it to meet customer needs, replatforming applications on new technology and the added significant cost of interoperating with other systems and national programs such as Spine, Referrals, EPS and MHRA medical device requirements, are all major barriers to new entrants, as well as incumbents. In addition, the costs of moving from one supplier to another are generally in

¹ <https://researchbriefings.files.parliament.uk/documents/POST-PN-0519/POST-PN-0519.pdf>

the hundreds of thousands of pounds, and often cause significant data-loss and disruption, which itself is a major disincentive for NHS organisations, despite the benefits of switching to a more modern or innovative product.

In fact, it could be argued that the ‘interoperability problem’ is just a symptom of a deeper [“wicked problem”](#) of capturing and exchanging health data. Such “wicked problems” need adaptive leaders who recognise that the solutions require new learning, innovation, an open-minded approach and a co-production mindset; no single technology, no single technical group, no single person can fix this alone.

HL7 FHIR[®]

HL7 FHIR[®] is the latest technology from the international HL7 non-profit organisation which has led global healthcare interoperability efforts for several decades. Unlike previous HL7 standards, HL7 FHIR[®] is a more complete Health Information Model and is open source.

Compared to previous HL7 versions, FHIR[®] is much easier for developers to work with and perform the critical, necessary and challenging curation process needed to turn raw ‘international FHIR[®] resources’, that make up the Health Information Model into locally (UK) interoperable local/national ‘profiles’. These UK-curated profiles are given the brand name “CareConnect”.

The [CareConnect curation work, led by INTEROPen](#), is [progressing well in the UK](#) and, when implemented will lead to significantly better data exchange for high-value data like medications, allergies, problems and immunisations between willing providers; accepting the status quo of different systems having different internal representations of clinical ideas, such as data models, workflow models and decision support rules. These CareConnect FHIR[®] profiles are focused and designed for data exchange between systems with different data models.

The raw international FHIR[®] resources adopt an ‘80/20’ approach to data modelling, concentrating on those data items that current suppliers **carry in their systems, are able to exchange, and for which there is demand for exchange**. FHIR[®] deliberately sets out not to ‘boil the ocean’ in terms of data modelling for less common data items, and was never designed to natively support the creation of full EPR systems, or the breadth of data modelling required for this more complex area.

Without any doubt, HL7 FHIR[®] represents a very significant advance on previous data exchange efforts and, for good reason, has significant backing from NHS England, NHS Digital, the other 4 country equivalents, and of course the INTEROPen community and its system suppliers.

For any policy-maker, CIO or CCIO, or indeed any supplier, it would be extremely surprising for HL7 FHIR[®] to be excluded from procurement documents, strategy or deployment.

openEHR

openEHR is a non-profit organisation dedicated to promoting the use of open standards and specifications but fundamentally re-imagines the way that health IT systems are built and deployed, following these base principles:

1. It must be possible to reduce the cost/resource burden to new market entrants of figuring out individually how to represent 'blood pressure', or 'medication' or 'allergy' as open and free technical artefacts, so that they can be shared and used internally for storage by existing and new implementers.
2. The clinical community should own and curate the openEHR clinical data models, or 'archetypes' (broadly akin to FHIR[®] resources/profiles), to improve quality and reduce unnecessary variability.

The openEHR community has developed a set of sophisticated tools to engage clinicians directly in building and reviewing these models. Note that similar to FHIR[®], although openEHR specifications and clinical models are open-source, the systems, applications and other related software stack need not be.

3. Applications should be separated from 'systems'. Typically users talk about 'hospital systems' or 'GP systems' to include the application, the data/data definitions, business rules and storage.

In the openEHR world the 'system' is only the data, the data definitions, business rules and workflow/decision support. It does not include the application itself. Those system definitions should be 'owned' and controlled by healthcare organisations, not by the vendor/implementer. Healthcare providers should be able to add or amend clinical content without any technical or contractual reliance on the system vendor. In contrast, FHIR[®] was primarily designed to allow the sharing of information, agnostic of who owns and controls the data. However that is only possible if/when the vendor is prepared or able to build the relevant FHIR[®] adaptors. Ultimately, in most cases, access to the datastore remains under their control.

As an example, suppose that a housing assessment model is agreed. Typically in a FHIR[®] environment, a profile would be developed and then system suppliers would be asked to read and write data conformant to that profile. They may need to change internal database structures and/or develop mappings to existing structures, and may not be able to fully support the profile. The scope, cost and timescales of those developments (if they happen at all) are controlled by the vendor and the IP of the internal data models remains with the vendor.

In contrast, in an openEHR-based system, any archetypes developed or used remain in the control of the provider.

The key aspect of openEHR technology underpinning this ability is the 'Clinical Data Repository' (CDR) which can be provided by any vendor, using their preferred technology (computer language, style of database etc.), as long as it is conformant with the application and database-agnostic [openEHR Reference Model specification](#); this reference model defines the basic rules that allow any openEHR clinical data model, or archetype, to be created.

New archetypes and templates can be deployed to the CDR immediately, and used to read, write and query the new data without any recourse to the CDR vendor, though additional work will be required to fully implement all the data flows associated with the new definition.

Of course, the application does also have to be written but is decoupled from the system, which allows different applications to work with the same CDR, or for applications to work with a CDR provided by a different vendor.

Data (and associated definitions) can be bulk transferred from one CDR to another without any lengthy or costly transform/mapping exercises (typically hours rather than months); often a major barrier to changing vendor. In addition the data is fully GDPR compliant - all patient data is available in a standard open format.

Developing an openEHR CDR is a non-trivial engineering exercise and which implies a much more significant commitment to a different system architecture, with an upfront cost. In many cases developers will opt to make use of existing third-party products rather than developing their own CDR.

4. This ecosystem must be designed in a way that allows new or adapted clinical data, workflow or decision support rule definitions, to be immediately applied to the EPR system without any re-engineering by the vendor.

All of the data recorded can be swapped between any CDR vendor system without long and costly data migration exercises. Application developers can have their apps work on different CDRs from different vendors. All of the interaction with the CDR is via the standard openEHR API, including data read/writes and querying. Apps do not work directly with data tables or equivalent NoSql constructs, but via the standard openEHR API which conforms to the Reference Model rules.

5. Because new clinical content and rules can be deployed in openEHR-based systems without any need for database changes or other technical modifications by the datastore provider, new clinical ideas and requirements can be deployed rapidly, and can also 'fail-fast'. Healthcare is complex, and seemingly innovative ideas/apps/clinical concepts frequently turn out to be flawed, have limited scope or are simply not viable commercially. We need to be able to build fast and fail fast with minimal cost, until truly great functionality emerges.

A good example is speciality registry apps: the data definitions are often shaped to support reporting needs but the apps themselves are increasingly used to support direct care. As they integrate further with other direct care systems there will be a need to replace many of the reporting-style definitions with those appropriate to direct care. e.g active problem lists where the diagnosis of diabetes is listed as a SNOMED CT coded entry, rather than boolean variables, such as “Diabetes, Yes / No”.

This process will be iterative as it must juggle the conflicting demands of reporting and direct care but, using openEHR, each iteration of data models can be defined very quickly using the tooling and immediately deployed once the models are built. No database changes, mappings or transforms are involved.

6. openEHR-based systems will, of course, always live in a wider ecosystem of non-openEHR vendor systems and will need to use ‘interoperability solutions’ like HL7v2 and FHIR[®] to share data with those systems. Therefore, almost every openEHR-based system is currently developing FHIR[®] interfaces to ensure they can support the patient’s journey between openEHR and non-openEHR systems.

What is different about openEHR is that those interfaces and transforms can be largely built once and shared between all openEHR vendors since they are based on the common openEHR archetypes; as explained in point 4 above, all instances of openEHR CDRs will natively understand any archetype without any further modification. It’s worth adding that as for FHIR[®], some degree of local templating/profiling may be required to support bespoke data requirements.

This “build the transform once” feature of openEHR opens up the possibility of sharing the burden of complex HL7 v2, v3, ITK or FHIR[®] integrations (e.g to Spine, SCR. IHR or a national ePrescribing system): a common software platform built on top of an openEHR compliant CDR can do the heavy lifting, allowing new market entrants to concentrate on developing truly innovative applications and functionality.

This is the approach suggested by the [Apperta Open Platform](#) definitions.

7. Where applications are clustered on an openEHR-based platform, they may not need data-exchange solutions like FHIR[®] because they are able to talk directly with other applications through the shared CDR and standard openEHR API. This, in essence, gives the level of tight data integration that would be seen in a single-vendor product but without the associated vendor lock-in. Of course, they will wish to have FHIR[®]-capacity to talk to the wider ecosystem, but they would almost certainly expect the platform to perform that role.

Is openEHR used in the UK or abroad?

openEHR-based systems/applications are growing in number and scope across the world with examples in the UK in London, Wales, Scotland, Plymouth, Taunton, Salford and Leeds, with others in the pipeline.

This is a technology which has been proven at scale, and is being used in Germany, France,, Australia, Netherlands, Slovenia, China, India, Norway, Finland, Moscow, Brazil, Spain, Italy, Malta, Portugal and the Philippines. These include full GP systems, hospital ePrescribing systems, regional and national EHRs with millions of patients, multi-hospital federated infection control systems and full paediatric EPR.

There are now multiple competing CDR vendors with some offering open-source products (openEHR is agnostic on business models). The number of patient records stored in openEHR CDRs is in the many millions, and the openEHR clinical content community is very active, and indeed has contributed both directly and indirectly to the design of many of the international FHIR[®] resources as well as local profiling efforts such as Care Connect.

A particular growth area for openEHR is the 'bi-modal system' approach described by [Gartner](#) where, rather than replacing existing established EPRs, the openEHR CDR is seen as a more agile partner; better able to address local demand for innovative, clinically-led app development. NHS Salford is adopting this approach with strong support from their EPR vendor.

New learning and open-minded answers

With the explanations above in mind, let's now return to the frequent questions: ***"Are NHS health providers and organisations expected to use HL7 FHIR[®]?"***

It is certain that any NHS organisation will see HL7 FHIR[®] as an important part of its technology portfolio and procurement requirements

"Why are some NHS health providers and organisation expecting to make use of openEHR as well as HL7 FHIR[®]?"

A number of NHS Trusts and other organisations are acquiring and deploying openEHR solutions, generally as an adjunct to their existing system portfolio, as a way of supporting innovation whilst minimising the incremental interoperability challenge of supporting 'feral' apps, and as a way of building a common data platform.

"Does openEHR 'do' interoperability?"

This question has generated some pretty heated debate amongst a number of INTEROPen board members; debate that is nevertheless welcome, necessary and reflects us living the community values of openness and transparency. Furthermore, it's the kind of debate we need to continue to have with the rest of the healthIT and INTEROPen community.

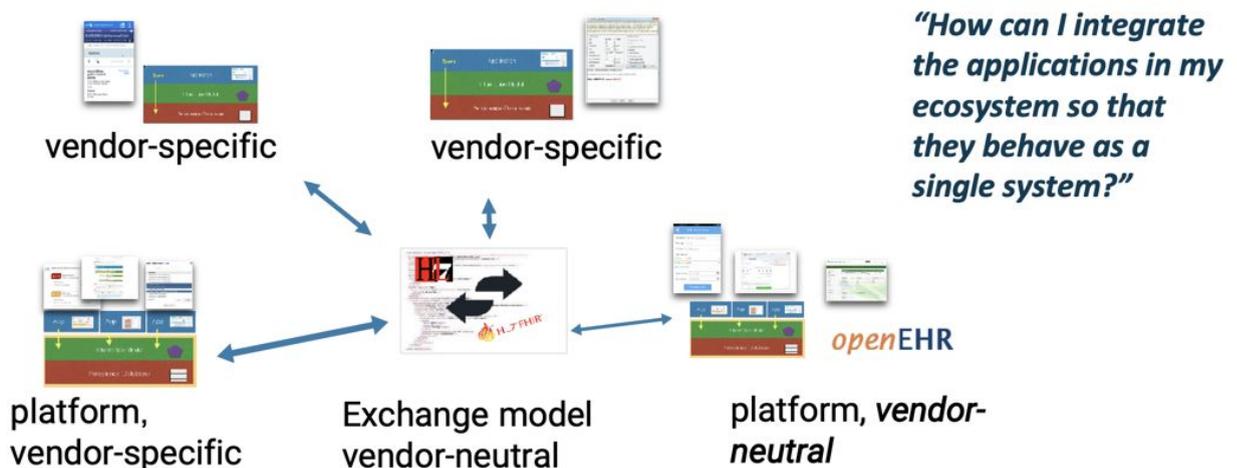
It is probably fair to say that although openEHR could technically be used as an ‘data exchange solution’, HL7 FHIR[®] accommodates the current needs and market state, is much less disruptive to existing vendors and has generated sufficient commitment from incumbent vendors to make it the sensible and practicable short and medium term approach.

Aside from the value of openEHR in the here and now, it can be seen from the information provided above that in the medium/long term, embracing openEHR affords us the chance to reduce the need for ‘interoperability’ between siloed applications by allowing those apps to share a common, tightly integrated but highly flexible common clinical datastore. It encourages a new breed of vendor, prepared to relinquish control of the data and data formats. One could say that openEHR indirectly addresses the ‘interoperability problem’, without needing to be regarded as a traditional ‘interoperability solution’.

This does challenge the current philosophy of tackling the data exchange needs between vendors and silos, and you will be forgiven many times over if you need to read this paper again and again for that “ah-ha” moment. The authors themselves have had to spend many hours, days, and sometimes years(!) getting their heads around this new way of thinking. Give yourself permission to be “open-minded” and “adaptive” in your approach to such “wicked problems”.

Time to play nicely together

As suggested earlier, where the overarching challenge is “how do we get this complex patchwork of applications to behave more like a single coherent system?”, then openEHR and FHIR[®] bring different approaches.



FHIR[®] pragmatically accepts the status quo of every system having its own internal definitions of clinical concepts, and limits its design scope to the exchange of commonly held, high-value data, with an expectation that, as the systems become more aligned, the process gets easier and the

cost benefit analysis improves. In the current world with many different vendor data models, FHIR[®] helps solve the data liquidity problem.

In contrast, openEHR offers a broader vision of a gradual, evolutionary but fundamental redesign of the way that healthIT systems are built; using commonly agreed and shared internal definitions, coupled with smart CDR datastores which can deploy these shared and evolving definitions without any new engineering effort. It sees the ‘interoperability problem’ as a symptom of a much deeper “wicked problem” of capturing and exchanging health data, and is attempting to reduce the need for interoperability solutions in the first place.

Although both groups are building clinical data models, these are often much more aligned than one might expect. In some cases there has simply been natural consensus, in others FHIR[®] has drawn on existing openEHR models, and vice-versa. In the case of allergies, the two communities worked directly together to develop a common model.

Arguably, the understandable focus on developing high-value FHIR[®] profiles risks ignoring the significant potential to allow more complex and detailed models to be first developed and clinically curated via openEHR tools and methodology; especially those that will be required to support the shared longitudinal record envisaged by LHCREs and the NHS Scotland Digital Platform.

A good exemplar of the potential for future joint working is the development by NHS Scotland, PRSB and key members of the Clinical ReSPECT team, of a set of openEHR models to digitise the new end of life ReSPECT form. The resulting template is now going into production as part of the longitudinal record underpinning the openEHR-based NHS Scotland Digital Platform. Further work will be done to develop a FHIR[®] wrapper around this template to allow for broader interoperability with non-openEHR systems but it will be much easier now that the detailed clinical requirements have been worked through.

This type of development, where technical teams support the clinicians and providers to develop the right solutions for their patients’ care is co-production in action; it’s not top down, but led by the service to develop useful, usable and used solutions. And it’s this mindset that helps us innovate by being open to new ways of working and new ways of thinking about the technologies we have and use.

Today, and for the short/medium term, we need both FHIR[®] and openEHR to deliver the integrated care we all want across the UK; but more importantly, our ability to deliver this vision depends on opening our minds and learning to respect, understand and support this brave new partnership.