The NHS England Technology Fund – Integration Options for Trusts

White Paper by HL7 UK

Abstract: The NHS England Technology Fund invites and supports NHS Trusts to make radical improvements in the level of their IT support, in order to provide better patient care. As well as the clinician-driven and patient-centric aspirations of the fund, the funding invitation sets out some technical and architectural themes, which Trusts should address in their proposals for funding.

Individual Trusts need to make hard choices – for instance, about what standards and APIs they intend to adopt, and the pace at which they will adopt them, in order to achieve local integration of their own IT applications, and to inter-operate with other local providers. They need to start making these choices rapidly – not just to satisfy NHS England that they will make good use of the investment, but also to set the right directions for themselves. Trusts know from experience that these are not easy choices to make.

This white paper by HL7 UK sets out some of the technical and architectural questions raised by the funding invitation, which NHS Trusts will need to address in the coming months. For these questions, we list some possible solution ingredients and discuss their relative merits.

If readers wish to make comments on this document, we would be grateful for them, and may be able to publish selected comments as an addendum. Please send your comments to webmaster@hl7.co.uk.

Contents

Executive Summary

In order to respond effectively to the NHS England Technology fund, each NHS Trust needs to set out its own **Integration Roadmap**, showing how it will use open APIs and standards to migrate from where it is now, towards the Integrated Digital Care Record (IDCR).

In this paper we set out some principles and ideas, based on our experience, which Trusts may wish to consider when designing their own roadmaps.

As set out in the Technology Fund, integration standards from HL7 can form a key part of such a roadmap.

- HL7 Version 2 is a well established standard that still has an important role to play. It is likely to be present in existing deployed products at Trusts as well as any new purchases. Its most appropriate use is for message based interactions: ordering, ePrescribing, database synchronization and so on. For replacing document based communications, or for building a complex health record, other HL7 standards are more appropriate.

- HL7 Version 3 aims higher than V2 with more sophisticated structures for complex interactions and EHRs. The implementation effort required may be too complex for local interactions. V3 is better suited to national level programmes with far reaching data interactions and well-resourced teams. Existing successful national V3 interfaces are now wrapped with APIs such as ITK Spine Mini Services, so Trusts can get the benefit of these without the effort.

- CDA, Clinical Document Architecture, is the world's leading electronic format for clinical documents. It is rapidly becoming the 'PDF' of healthcare, and is well supported already by ITK and MIM/DMS. We see CDA as a vital part of an integrated workflow within IDCR. CDA is a standardised document format for healthcare, deployed or planned in over 30 countries. It was recently the subject of several successful and rapidly developed projects funded by ISCF.

- FHIR (pronounced 'fire'), Fast Health Interoperability Resources, is a new free and open standard API for medical applications. FHIR offers a quick and easy way to get systems connected using off the shelf web standards, familiar to even junior developers. In late 2013, it will become an HL7 Draft Standard for Trial Use (DSTU). Right now, it would make a great basis for a newly constructed interface for in-house systems. However commercial support is lacking as yet, so it can't be used as part of a bought in deployment. FHIR is designed to encompass the capabilities of all of the other standards, keeping it simple and focussing on the 80% that is really needed. It's a likely long term single platform standard that all the others can migrate to.

Technical and Architectural Issues Raised by the Fund

This section summarises some of the technical and architectural issues raised by the Technology Fund, illustrated by selective quotations from 'Safer Hospitals, Safer Wards: Achieving an Integrated Digital Care Record'. In what follows, we shall refer to that document as 'SHSW'.

Sections 1-3 of SHSW set out the broad goals of the fund, which we only briefly summarise here, assuming that readers are familiar with them. Section 4, on Architecture and Standards, sets out some of the key technical questions which arise along the route to those goals, and it is those questions which we address in this paper.

Section 1-3 Summary Points

Sections 1-3 of SHSW set out the aims of the fund, which is to help Trusts towards the vision of a '*fully integrated digital patient record across all care settings by 2018*'; where the meaning of 'integrated' is that '*vital patient related information and clinical decision and support tools can be viewed by an authorised user in a joined up manner in any single instance*'

This theme of integration of information from different IT systems runs through SHSW – if not as the initial objective ('safe digital record keeping') but at least as the direction of travel, to be achieved by 2018. The proposed solutions have a number of aspects, most notably 'Open APIs' and 'standards-based interoperability'. These themes are addressed in section 4 of SHSW.

Section 4: Architecture and Standards

This section of SHSW '*outlines the different approaches for developing the technology stack of an Integrated Digital Care Record (IDCR) and the underpinning principles, technical capabilities and key standards that local teams should consider when implementing safe digital care records*'.

Having recognised a range of possible IT architectures from 'single solution' through to 'best of breed', SHSW then addresses a list of 7 topics which are relevant for any architecture:

1. Patient identity
2. Digital Data Capture
3. Storage and Indexing
4. Open APIs
5. Patient access to records
6. Information governance
7. Standards

Probably every Trust in England needs to make significant progress in all of these areas. Not wishing to under-state that challenge, for the remainder of this paper we shall address just two of areas: (4) Open APIs and (7) Standards. Choosing an appropriate direction here will influence the success of Trust IT projects for many years to come, as well as affecting the success of any bid for funding. It should be noted additionally that at a technology level, item 1 can be facilitated in large part by synchronizing patient NHS numbers using Open APIs and Standards.

Item 4: Open APIs

SHSW sees an important role for Open APIs in any Trust's IT architectures, stating that '*When purchasing new 'best of breed' systems, the flow of information in and out of the system will be enabled through Open Application Programming Interfaces (APIs). Specification of Open APIs should form a part of the procurement criteria for Trusts purchasing digital systems both for core systems and for 'integration layers' such as portals/integration engines*'.

SHSW states further that '*Open APIs allow modules and systems to integrate seamlessly with one another in a standard way*.'

What does that mean in practice? In our view, Open APIs are a worthwhile step towards local integration of applications but they are far from being the whole story. For that, you need to progress from Open APIs to s*tandard* Open APIs.

If all Trusts were to require open, published and documented APIs on the systems they procure, subsequent integration work could in principle be done by others – including the Trust's own IT staff, opening the work up to competition and helping to reduce costs.

Even with this you may not get the '*seamless integration'* envisaged by SHSW, because different open APIs may be completely incompatible. Integrating system A with system B, requires an in-depth understanding of both APIs, and often requires considerable effort to bridge the gap between them. When a Trust has several incompatible systems, the number of ways these have to be connected can rise rapidly - and become an unacceptable workload.

Draft supplementary guidance issued by NHS England explains what is expected from Open APIs.

Essentially the requirements are:

- Read and write records, including patient demographics and the full clinical record
- Search for records, by criteria
- Accessing key system functionality, e.g. launch application with a certain patient showing.

The remainder of features specified are more qualitative, concerning privacy and security, scalability and completeness of documentation - which are of course prerequisites for any product in this area.

Consequently, this boils down to some fairly simple actions being supported: Read data (by identifying number - 'id'), Write/update data (by id), Search (by parameters) and Launch Feature (by id - perhaps requiring a feature or user interface component identifier, plus a data item id).

With this fairly limited feature list it seems easy to accomplish connecting to such an API, standard or otherwise. An API that has half a dozen or so parameterised capabilities typically goes a long way to getting applications to interoperate in a useful fashion. The hard part is almost hidden, and is in the definition of the data format that is to be read, written or retrieved by a search.

We consider that, while a standard way to access application actions is nice to have, the real benefit to standardising is in the data formats themselves. This is a case where the devil is in the detail. The complexity of a request to "read patient where id=#####" is trivial compared to the 20 or so

structured data items retrieved for, for example, each medication authorisation item applicable. When you add in full demographics, lab results (and requests), dispense items and so on, the data definitions become significant. Even details such as how dates and times, dose formats, normal test ranges etc. are represented can require much work.

HL7 V2 has well established standardised ways to read, write and search for data between conforming systems, covering all these use cases. More importantly, it underpins this with a standard well documented data representation for these, across all application types.

The forthcoming FHIR standard (see page 10) updates HL7 V2 with a web standards based method for reading and writing data directly between applications. FHIR uses HTTP-based 'RESTFul' style APIs that that are now industry standard (as used by Amazon, Google and PayPal). FHIR has been designed as a free open standard healthcare API from the start. Its simplicity is illustrated by an example of how you search for demographics of patients called 'Smith', simply by accessing an URL in a format such as: https://localservername/patient/search?name=Smith

(This assumes internet standard access controls are also in place of course, as commonly used in banking and other secure situations.)

FHIR Is not yet built into commercial systems, but would make an ideal start for a web standards based API into existing systems, where no API is yet exposed, and is something to ask vendors to support.

Item 7: Standards

This section of SHSW states that standards are required so that '*the meaning of each item can be established consistently both by users that may use the information to drive decision support or by automated workflows.*'; but also recognises that '*Semantic harmonisation is a journey. It will take time for organisations to achieve it. Organisations need to start the journey of digitisation without trying to achieve perfection before starting to share clinical data across specialties*'.

It then lists some of the standards which may help along this journey, under a number of headings:

Key standards - Semantic:

- SNOMED CT
- NHS classifications:
    - OPCS4
    - ICD10
- dm+d
- UCUM
- NHS Data Dictionary
- ODS – (formerly NACS)
- ISB
- PRSB
- Academy of Medical Royal Colleges Record Keeping Standards

Messaging:

- CDA
- XML
- MIM
- DMS
- HL7 V2
- HL7 V3

Transport:

- Transport Layer Security (TLS) – Approved Cryptographic Standard

How are Trusts to interpret these lists when formulating their own roadmap, for the journey towards semantic harmonisation and IDCR?

It would take too much time and in-house skills acquisition to understand and apply all the various standards listed. Even to select a key subset to start from would require some consideration. As a starting point, we present below some observations on the relevant HL7 standards, based on our experience.

Solution Components – History and Experience

HL7 Version 2

HL7 V2 is a set of standards for healthcare messaging which is more than twenty years old and continues to be very successful in the market, both in the UK and internationally. It has coverage of many healthcare domains, and there are V2 interfaces to many suppliers' systems; for instance, ADT messages to and from Patient Administration Systems. V2 exchanges are a common part of the IT landscape in most Trusts.

There is a profile of HL7 V2 defined as part of the NHS Interoperability Toolkit (ITK). Version 2 predates XML, but an XML version has existed for a long while and is used by ITK.

V2 is good for most of the bread and butter flows within the Trust. Document-centric data, such as discharge or incoming referral letters is a better fit for the CDA standard (q.v.), but most other intra-Trust requirements can be met with V2 interfaces. Laboratory order requesting and results, medication and pharmacy messaging, including ePrescribing, authorising, dispensing and administering are all well covered, as are demographic links from PAS to other systems.

We see V2 as a cornerstone of Trust level integration.

Version 2 continues to be developed and modernised by HL7, with version 2.8.1 being proposed as a new standard in July 2013.

HL7 V3

HL7 Version 3 was defined by HL7 International to modernise and extend the capabilities of Version 2. However V3 hasn't displaced V2 in most hospital settings, because V2 does a good job within its traditional role. V3 is intended for more sophisticated interoperability tasks where the complexity begins to overwhelm what is feasible in HL7 V2. V3 is suitable for extracting an entire health record or full medication history, and as such is used in the GP2GP system that moves the complete structured patient record between all the GP systems in the UK. This tends to be overkill in the Trust environment, where more focussed messages to carry a set of lab tests, or a medication order, are typical. These messages flow to build up the EHR, but the full EHR itself is not normally represented in HL7 Version 2.

HL7 Version 3 has been used to implement several successful national systems including GP2GP and eTP (electronic transmission of primary care prescriptions), as well as the NHS number tracing service (PDS), which helps to address another key part of SHSW (Patient Identity, SHSW section 4.2.1). In the latter case the V3 messages are now wrapped by the ITK Spine Mini Services, eliminating the need for Trusts to deal with V3 directly.

HL7 CDA

In our view the Clinical Document Architecture (CDA) will be an important component of any Trusts' interoperability roadmap, because exchange of clinical documents is a vital requirement, and there exists an established and developing national direction for implementing CDA.

CDA is a specialisation and simplification of HL7 Version 3 which has been created to support attested clinical documents. These can be a mixture of natural language text, displayable on any device with a browser, and structured coded data (e.g. a set of medications). CDA places the text at the forefront, which an authoring clinician attests as a complete and consistent statement, and the coded data may have a subsidiary role. CDA requires that any coded data must be fully consistent with the text visible to the clinicians, and it allows the codes to be linked to a particular phrase in the narrative.

This mix of text and optional coding makes CDA easy to implement, and yet it can scale up to the most sophisticated use cases.

There is a well-defined 'maturity path' for CDA implementations (parallel to the Trust maturity levels proposed in SHSW), which Trusts may follow as part of their interoperability roadmap:

- **CDA Level 1**: This includes only CDA header information (such as information about the author, patient, record keeper, message sender etc.) in a structured XML form, and may contain the body of the clinical document in any convenient form – even a PDF. This is a limited use of CDA and does not allow software to access the individual information items with the body text. However it is a still a useful way to index and move human readable clinical information around, and is a first step towards interoperability. In most cases though a level 2 implementation is preferable and achievable.

- **CDA Level 2**: This includes the same informative header as in level 1; but the body of the document is conveyed in a set of sections and sub-sections, each of which may contain a

heading and text or HTML tables of information. These help organize the document into standardised and consistent formats. Headings may be coded to allow software to recognise the different categories (see RCP standard headings, below). A simple and free 'stylesheet' allows the sections and data to be displayed on anything that is capable of showing HTML (from desktops to tablets).

- **CDA Level 3**: This uses the same CDA header as levels 1 and 2; but in CDA Level 3 clinical codes are added as an extra level of detail for the same text. Level 3 CDA is still fully displayable and readable as HTML text, but codes are also present for machines to process. To a human, level 3 and level 2 look identical in use. Not every piece of text needs to be coded - in some cases certain sections are plain text and others entirely coded. In a discharge summary, a written report may be accompanied by coded and structured medication, using dm+d SNOMED codes with computable information for dosing and timing. The GP may be able to incorporate some of this into their patient record. CDA supports any and all code systems, including SNOMED CT, and UCUM units (e.g. for lab results).

At level 3, detailed definitions of what is expected for each document are normally created, describing the rules for each section that the software will create. These 'templates', based on analysed requirements, define the structure, and the actual content of the CDA is filled in according to these rules, when the document is created and populated with data. Templates can be reused to create different CDA communication types. For instance, in the ITK there are CDA medication templates which define standardised representations of medication information, and that can be plugged into any CDA implementation.

There are now successful UK deployments of all three levels of ITK CDA. The levels form an orderly progression, with digestible increments in technical complexity between each level.

CDA has achieved a great deal of acceptance in the market. In the US, government requirements for 'Meaningful Use' have driven a high level of supplier interest in CDA, and a particular specialisation known as Consolidated CDA. In Europe the epSOS cross border medication summary is being piloted in almost every nation, and there are CDA projects running in South America, Japan, and Australia, to name a few. (See the XDS deployment map below).

The Interoperability Toolkit (ITK) has adopted many of the CDA definitions coming out of the National Program, and has encouraged suppliers and Trusts to implement them, largely through the Information Sharing Challenge Fund (ISCF). This has led (so far) to something of the order of 30 successful implementations across the UK, all completed within the challenging 6-month timescale of the ISCF. Trusts may wish to examine some of the case studies of ISCF CDA Implementations, to learn about the practicalities of implementing CDA.

There is a further strong reason to consider ITK CDA as a next step in each Trust's interoperability roadmap, concerning the Technology Fund's goals of clinical engagement. The Professional Records Standards Body (PRSB) has recently been set up to be '*the first point of call for professionals, professional organisations, service providers, commissioners, policy makers and system suppliers for expertise and all matters relating to care records*'. PRSB have published new standards in July 2013, with the endorsement of the academy of Royal Medical Colleges and over 50 clinical professional

bodies, and drawing on the earlier RCP Headings for clinical communications. HSCIC are developing technical CDA specifications to match these standards, for first release in October 2013. Those CDAs will be the pre-eminent choice for clinically endorsed, interoperable exchanges of clinical documents.

HL7 FHIR

Fast Healthcare Integration Resources (FHIR, pronounced 'fire') is a new healthcare standard from HL7 which is not in the list of standards cited in SHSW. However, in our opinion it merits serious consideration for the interoperability roadmap of any Trust.

FHIR is an international HL7 initiative (not US-dominated) which has been developed rapidly over the past two years. The current version of FHIR has been 'frozen' in July 2013 for HL7 International Ballot in September as a 'Draft Standard for Trial Use' (DSTU). Further fine tuning changes to the DSTU may be balloted in January 2014. For the period of the Technology Fund, there will be a stable version of FHIR available for implementation.

FHIR has been developed by HL7 to be a modern, easy to use standard. This has made it much more developer-friendly and easy to implement rapidly, as several successful 'FHIR Connectathon' events have demonstrated. They have shown that new developers can learn FHIR and build a working FHIR interface, within one or two days. In some ways FHIR is like HL7 V2 but updated for modern web standards, but it also extends to what was previously only easy to achieve in V3.

Some of the key characteristics of FHIR are:

- **Focus on implementers**
  The FHIR specification is written for a target audience: implementers. There are publicly available test servers. Starter APIs are published with the specification (for Java, C#, etc.). Lots of examples are available.
- Keep **common scenarios** simple
  It should be possible to 'figure out' FHIR, over a weekend. Easy to get started, then grow into the specification for more complex scenarios. Don't try to cover all: inclusion of content in core specification is based on '80/20' rule.
- Leverage **existing technologies**
  Web and mobile development technologies exist and are highly successful. FHIR reuses these: XML, JSON, REST, HTTPS, OAuth.
- Provide **human readability**
  Part of the reason for CDA's success is that everything is human readable as a minimum.
- Make content **freely available**
  FHIR is a free, open API, with free reference implementations:
  http://hl7.org/implement/standards/fhir/

With these advantages, FHIR also does not ignore bridges to other standards. In its first release FHIR is being mapped back to other standards V2, V3 and CDA.

Development and trialling of the FHIR standard has been very rapid in comparison to most international standards, and has been based on implementation experience at all stages. The level of

healthcare IT supplier interest in FHIR is very high, particularly amongst suppliers of integration tools. FHIR is being seen as a key emerging technology, with potential early adopters being HSCIC and IHE.

For all these reasons, we can expect FHIR to be an important standard in the coming months, and years, perhaps the dominant standard leading up to the 2018 target date for IDCR.

With that said, FHIR is not yet available in commercial products. Right now FHIR is best suited as a technology for implementing a bespoke integration between Trust systems, where the alternative would be a home grown and designed API. Using FHIR is easy, easier than creating your own API, which is the common reason many avoid standards – the advantages are known, but the cost of entry is sometimes too high. If a Trust has this sort of requirement, they should look closely at FHIR. It prides itself on being able to be understood in a weekend, and to be up and running in a basic fashion in a matter of days.

Because FHIR is so much simpler and more developer-friendly than the other healthcare information standards, we would recommend that those Trusts with a strong in-house IT capability should soon get some hands-on developer experience of FHIR. This could lead to useful specific applications of FHIR in the near term; but more importantly, it will help the Trust better assess the possible future roles for FHIR in its own interoperability road map, on the basis of direct experience.

HL7 UK offers training in FHIR, has it as a regular focus of its Technical Committee Meetings, and our members are closely involved in the take up of the standard.

IHE XDS

XDS is a profile for sharing of clinical documents in a repository, which has been defined by Integrating the Healthcare Enterprise (IHE), and is widely deployed internationally, although XDS currently has only limited deployment in the UK. A map of XDS deployments, many also using CDA, is linked from this page http://hl7.org.uk/version3group/cda.asp.

Sharing of clinical documents in repositories has an important role to play in achieving the goals of SHSW and IDCR, alongside system-to-system interfaces, in a number of respects:

- Giving a team of clinicians, or carers from a range of specialities including social care, shared access to a set of clinical documents, is a very effective way for them to share information as they need it (e.g. from mobile devices), rather than when some application thinks they need it.
- With appropriate controls, repositories of documents may be made accessible to patients and their relatives and carers, moving towards the objectives of patient involvement and empowerment.

Document repositories therefore have an important role to play in the interoperability road map; and the technology of choice to implement them is XDS.

An XDS repository is a collection of clinical documents, together with some indexing ‘metadata’ which enables users to search the repository for documents of interest to themselves – for instance,

about a particular patient. Correctly populating the XDS metadata is key to making an XDS repository useful.

There is a natural synergy between XDS and CDA. If the documents in an XDS repository are CDA documents, then the header information in each CDA (and sometimes, information in the CDA body) can be used to populate the XDS metadata which enables searches for documents.

HSCIC, HL7 UK, and IHE UK have worked together to ensure that CDA header data and XDS metadata are aligned, so it is clearly defined how to populate XDS metadata from CDAs. There have been at least two successful ISCF projects in which CDA and XDS have been combined using these principles.

SNOMED CT

The encoding of information about medical conditions, treatments, medications and so on has always been a core component of clinical interoperability, and remains a highly challenging area. Local codesets are often in use, and difficulties arise when these need to be mapped or translated into other local codesets, or standard codesets.

Amongst standard clinical codesets, the one with the greatest coverage, precision and international commitment is undoubtedly SNOMED CT, as listed in SHSW.

SNOMED is not without complexity however and it is therefore our view that the progression of any Trust towards implementation of SNOMED CT coding should be in modest incremental steps, with strong central guidance from NHS England and HSCIC about which domains have the most mature coverage and offer the greatest benefits from interoperability.

The main content standards, including all those from HL7, already provide the technical means to carry all uses of SNOMED CT codes. dm+d is a part of SNOMED and so also covered, as are other NHS classifications. The UCUM system is fully compatible with HL7 standards.

A Clinical Integration Roadmap for a Trust

In the previous sections we have given evaluations, based on our experience, of some of the standards components available to Trusts in their journey to the 2018 IDCR. Each Trust will need to evaluate some of these components, in the light of their current context and capabilities, and we hope the comments above provide initial guidance on the components currently available on the shopping list; more detailed guidance is available from HL7 UK and its members.

To formulate and implement a viable local integration strategy, Trusts need more than a standards shopping list. They need a roadmap – a coherent set of standards choices and migrations, enabling them to move from where they are now, to where they need to be to meet the goal of an IDCR in 2018 – with the minimum of cost and difficulty, and with the maximum of openness, autonomy, flexibility to meet local clinical drivers, and competition amongst their suppliers.

The actual roadmap will differ for every Trust according to their local context and capabilities; it will be complex and multi-faceted, and have individually-set timescales.

We set down here some principles which Trusts may wish to consider, when designing their own interoperability roadmaps.

1. **Any roadmap will save you time and money, compared with no roadmap at all**: This principle borders on the obvious, but is often not heeded in practice: if the steps you take are just a series of tactical moves, with no guidance or end point, then you will waste time and money (and more important, lose clinical commitment) in unnecessary re-implementation and re-work. The earlier you draw up a coherent roadmap the better; so devote ring-fenced effort to drawing one up in the first three months of the Technology Fund, and subjecting it to wide critical review – and then annually to revising it in the light of the progress you have made.

2. **Base the roadmap on a thorough and realistic assessment of where you are now**: This means not only a clear analysis of your IT estate – in diagrams, spreadsheets and other forms – but crucially, the clinicians' view as well; where do these IT systems serve them well, and which of their needs are not met? Build in the clinicians' view to the roadmap from the start, in order to gain and keep their commitment to the changes which will occur. A third viewpoint, which spans the technical and clinical view (and will be crucial in your migration to the universal use of NHS number as an identifier) is the **data quality view** – what information resides in which IT systems, and how good is the quality and consistency of that information? Many change programs have failed, through failure to appreciate the seriousness of initial data quality issues.

3. **Plan Modest Incremental Steps in Interoperability:** The history of healthcare IT is littered with 'great leaps forward' which centred on huge technology-led investments and led nowhere. Instead, engage with your clinical staff in the art of the possible – given where you are now, your intended direction of travel, and the tools and standards available today, what modest and feasible next steps in integration will provide clinical benefit and build the confidence for further improvements? For instance, the three levels of CDA provide a graded maturity path for electronic clinical communications, and are applicable across many clinical specialties.

4. **Regularly assess the maturity, take-up and complexity of available standards:** In the alphabet soup of healthcare standards, every standard has its proponents; but in choosing your own route, you will need to apply hard-headed tests of maturity, deployment, and concensus with the other care providers you will interwork with. For instance, the new HL7 FHIR standard shows considerable promise in terms of simplicity and ease of implementation; but as of today, it is lacking in supplier support and deployed implementations. This may be a fair summary of the situation today; but given the rapid worldwide development of FHIR, will it still apply in 9 months' time? Actively assess FHIR now, so you are ready to use it when you need it.

5. **Define your direction of travel, and then monitor progress against it:** The preferred direction of travel is from diversity to greater simplicity – from a plethora of different representations for healthcare messages, documents and APIs, towards a smaller set which

maximises the ease of interfacing between different applications, to truly integrate the IDCR. This smaller set might not be one all-encompassing standard; but it may be. How should you measure the diversity of your APIs, and measure your progress in reducing it? What are the migration steps you need to take?

These are some of the key considerations in designing an interoperability roadmap. It is very worthwhile to strive to get it right, and then to apply the roadmap in all your IT procurements. The difference could be a large part of your total IT budget, and a huge difference in its clinical effectiveness, over many years.

In drawing up and reviewing the roadmap, seek a diversity of opinions. Do not rely solely on your own in-house IT staff, or the word of your major IT suppliers. HL7 UK have many members who are willing and able to give you independent and expert advice, to ensure that your own interoperability roadmap is soundly based, practical, and best meets your clinical needs. Some of these consultants are listed at http://www.hl7.org.uk/techfundadvisors.asp. HL7 UK members can also use the discussion lists and attend free technical meetings.

The Role of NHS VistA

One of the more controversial aspects of the NHS England Technology Fund is its intention to develop 'NHS VistA', and make it available to Trusts. Expectations as to the delivery dates of NHS VistA are being sensibly lowered however – '*we do not envisage the roll out of an Open Source IDCR solution at scale or pace within the life cycle of this Fund.'*

The view of HL7 UK is that any system deployed into Trusts, including a variant VistA, would need to work with the existing pattern of common open APIs and standards, as described elsewhere in the document, rather than dictate any new ones. So the current uncertainty about the nature and timescales for NHS VistA need not delay the essential work of setting out each Trust's own interoperability roadmap.

Conclusions

This paper has surveyed some of the issues now facing NHS Trusts in their choice of strategy for achieving the 2018 target of an Integrated Clinical Digital Record (IDCR). We note that to have any prospect of meeting that target, the use of clinical information standards is essential - if only to restrict the number of individual system-to-system interfaces which would otherwise need to be built.

We have surveyed some of the standards available today, summarising the experience in the UK of practically implementing those standards. We noted that to choose its strategy, a Trust needs more than just a shopping list of standards to apply; it needs a coherent road-map, which recognises how the different standards fit together (or do not) and gives a feasible incremental migration path from the Trust's current state to the future IDCR.

Appendix A: CDA Implementation Strategies

CDA implementation will be new to many Trusts, and while support is increasingly available in off the shelf software from suppliers, there are a range of techniques to make local implementations easy.

CDA is a capable and far reaching standard, and consequently allows implementations of some complexity – or something quite simple.

A number approaches for both reading and writing CDA documents have been developed. These techniques include:

- **'Filling in the blanks' implementations**: Here, an existing example of the CDA XML is used as a skeleton, and the task of the developer is only to build software to 'fill in the blanks' of this empty XML structure. This is then standard technical task that is not unique to CDA and is viable for any up to date developer.

- **CDA APIs**: These are APIs in some language such as C# or Java, which allow the developer to construct a CDA instance using simple calls to the API – for instance, using 'setter and getter' methods to access the patient's date of birth, gender and so on. Then code within the API takes over the job of reading or writing the CDA XML format for you.

- **Green CDA**: In this approach a different XML structure is created, which is simpler, smaller and flatter XML than what is allowed by full CDA. This structure can be made specific to the particular Trust application, and can use names that programmers are familiar with. This is then easier for a non-CDA expect to work with. As part of the method, this new structure can them be converted to or from full CDA XML by reliable, automatically generated transforms. The developer need only interface his application to the "green" CDA, a standard XML interfacing task requiring no CDA specific information, and the generated transforms then support exchanges of full CDA with other applications.

In each of these approaches, developers are shielded from the more complicated parts of the CDA specification. Within the ISCF there are successful, accredited examples of all three of these approaches to simplified CDA implementation.

As a balancing factor, these methods do not always allow access to every possible attribute of CDA. They each support read and write access to a form of '80/20 CDA', which supports 80% of application requirements with 20% of the technical complexity. Since that is the CDA subset which normally used, that is sufficient for most interoperability tasks. It is always possible to extend the methods to add extra parts of CDA, where necessary.