

Interoperability: Banking did it... how hard can it be?

by **Shawn Vincent, Vice President, R&D, MD Practice Software LP**

Sep 13 2012

I have been in the business of designing and interconnecting electronic medical records (EMR) systems for a few years now. My particular passion is connecting EMRs to each other and to other systems. For years, there has been a very powerful dream within the healthcare IT (HIT) industry that all of the various products from all of the various vendors will one day be able talk to one another. They will combine into a single system – where all data about a patient is available at any point of care. This dream, typically, is enabled through “healthcare standards.”

Many people who speculate about this dream typically neglect to describe the unique challenges that need to be overcome. I will attempt to illuminate some of these challenges here – and add some thoughts about what we can do about it.

There are a number of myths regarding healthcare IT. Let’s talk about one of those myths today, explain why it’s false and the problems that it causes.

The myth is: “Banking did it... how hard can it be?”

It’s not always “banking,” of course. Folks also tend to mention another industry – travel – which has managed to connect together systems from various vendors across the industry, and which is able to achieve a high level of interoperability, as another prime example of the myth in action. “Standards” are typically a component of this argument as well.

But let’s concentrate on banking. You can walk outside right now, to any bank machine, for any bank, slip in your universal banking card and access your money. Somehow, all of the world’s banks have managed to come up with standards and connected systems that make all of this work. Sure, there are occasional glitches – but, by and large, it all works really well.

Imagine if HIT could achieve this level of interoperability. You could go to any doctor, clinic, physiotherapist, nursing home, etc. – and get immediate access to your full medical record, just like banking! You could even connect and see your own health record on the internet, using secure web browsing, just like banking!

At this point in the conversation, the myth continues along these lines: “Health care is 20 years behind the rest of the world. If vendors could just agree to adopt some universal standard for healthcare data, we could bring HIT kicking and screaming into the twenty-first century!”

There are many reasons for HIT’s lag when compared with other industries. Politics and costs play a big role here, but we’re not going to talk about these just now. Instead, let’s concentrate on the technical problem itself. I will argue that achieving banking- or travel-level interoperability within the healthcare system is much more complex than what is happening in those other industries. “Just add standards” isn’t as easy as it sounds!

In fact, I will argue that the single-minded devotion to a “universal healthcare standard” (sometimes, you’ll hear “USB for health care”) is harmful to the achievement of this goal.

First, to set things straight, standards for transmitting healthcare data have been in the works for a long time. HL7v3, a truly universal healthcare standard, has been in the works since 1995 – that’s 16 years ago. SNOMED CT, another popular universal standard, has been in the works since 2002 – and is based on earlier coding systems, which had been in development since the early 1980s.

Adoption of these standards, even today, is limited. There are a few high-profile sites that are using them – but not in ways that would enable them to talk to one another other than trivially. Most uses of HL7v3 and SNOMED CT incorporate dialects of the standards – which means that the standard itself is modified for use within a particular facility or region.

“Modify” a standard? That’s crazy! Why would a “standard” be modified? There are many reasons, but two stand out: 1) They’re too complex to implement fully, leading to people sub-setting the standards; and 2) Translation to and from legacy systems is typically too costly (and sometimes impossible), leading to people extending or modifying the standards.

Because the standards are modified when used, a message that’s produced in one facility or region cannot be transmitted to another facility or region without first being translated. Note that if the two regions were not using the standard, you could still transmit messages between them if you were willing to translate the messages. Note too that the world with the standards isn’t actually much different than the world without the standards. You still need the translation.

So, why, if the standards exist, is the healthcare industry so far behind banking and travel? What’s the difference? Why are people modifying the standards? What makes this translation stuff so hard?

Once you’re in the trenches, it turns out that the problem that banking solves is not in the same league as the problem that health care has to solve to achieve true interoperability.

Why? The key reason is this: In most industries (like banking), the business constrains the data model. In health care, the data model cannot be constrained.

Let's take a good look at banking. Bank accounts are relatively simple data structures – they have only a few data attributes: a customer ID, an account balance... maybe 50 attributes at the end of the day. They can be simple – in fact, they have to be simple. If they got too complex, the bank couldn't function as a business – bank processes are built entirely around this data structure; or, rather, the data structure is built entirely around the banking business.

The complexity of a bank account's data structure is constrained by the complexity that the business needs and wants to deal with.

Think about the alternative. If every possible aspect of a bank account had to be modelled, you'd need to incorporate people's budgets, their lifestyle choices, their stock preferences, maybe even their sister's names! You'd need to record stuff you never even thought about until the customer came up with it. You'd need to train all of your staff on how to enter all of this data. You'd need to ensure that every staff member is using the system correctly and consistently. The result would be skyrocketing costs. Banks are cost-driven, so you tend to keep the data model really constrained.

Take another look at the hypothetical banking situation I've outlined in the previous paragraph: This is exactly the situation that exists in health care. You need to model pretty much everything, even the patient's sister's names! Health care models the real world: biology, medical science and treatments. These things change over time; new stuff is added constantly. The world has infinite variety and complexity – you'd need to model lots of really complicated stuff!

Further, you don't have the luxury of having a choice to simplify it! You can't say, "Don't bother recording some detail of a patient's health because our computers can't cope..." – the computers ALWAYS need to cope.

Let's look at some real-life examples of the complexity of healthcare data. (Disclaimer: I am not a doctor; I just hang out around them a lot in my job.)

Example 1: One of my favourite examples is from SNOMED CT. It demonstrates that medical data can be really weird.

It starts with psoriasis, a symptom (flaky skin). This symptom is associated with a bone condition known as psoriasis with arthropathy. This condition is often found in juveniles, where it is known as juvenile psoriatic arthritis. Psoriatic arthritis sometimes doesn't present with the skin symptom, which leads us to juvenile psoriatic arthritis without psoriasis. In technical terms: Wow.ⁱ

This is the sort of thing that gives software engineers like me nightmares. How do you cope with complexity like this? SNOMED CT models this – and even did a good job – but it's complexity like this that makes standards like SNOMED CT very difficult to implement well.

Example 2: My experience in writing computer physician order entry (CPOE) systems has taught me that different clinical stakeholders think at different levels of detail. Let's take a look at doctors versus pharmacists in something that they both deal with daily: prescriptions.

Doctors want to think about doses and some kind of reference to a drug (think "20mg acetaminophen" here). The doc doesn't generally care about the pill size/strength or the brand name. Note here too that there are no standards for generic drugs.

Pharmacists, on the other hand, need to think about both strength and manufacturer (think "2 10mg tablets of Tylenol" here). The strength and manufacturer are core concepts to the pharmacist. Note here too that the pharmacist's version of a prescription is formulated at a lower level of detail than the doctor's preferred view.

Why the difference? The pharmacist needs to manage the pharmacy inventory, and dispense the correct number of pills. The doctor wants to treat the patient.

Most computer-based drug systems are built for pharmacists, because they've been computerized longer than CPOE systems. The clinical standards that apply to medications were built for pharmacists. There are no standards in place here for what the doctor wants to say – and even if there was, how do you translate that into what the pharmacist wants? In the paper-based world, there are humans who do the translation. In a fully computerized world, the problem space gets complex really fast!

Conclusion and suggestions

These two examples represent just the tip of the iceberg. I could go on for hours (try me sometime!) Suffice to say, the data found in health care is complex and weird – and traditional computer systems aren't great at coping with it. As a result, different systems from different vendors take different approaches, and when you try to translate data from one to another, all sorts of hijinks ensue.

As a result, it should come as no great surprise that health care has a much more difficult "problem" to solve than banking ever did. "Just add standards" isn't an approach that will work.

Is it possible that, with a sufficiently refined standard, the standard will bring about the dream? Maybe – but in practice, the industry is finding that once you have a standard that handles all of the subtleties and complexities of real-world health data, the standard becomes too complex to actually implement when using real technologies and real programmers in real situations. Throw legacy systems and historic data into the mix and what you get is a real mess.

Meanwhile, the myth of "just add standards and catch up to banking" continues – as does an obsessive, almost manic dedication to healthcare standards in this industry – to the exclusion of all else.

But this myth, coupled with the fanatic dedication to standards, is particularly dangerous, because it distracts us from seriously looking at approaches that are based on real-world challenges.

So, is the problem unsolvable? Of course not! However, we will need to focus on strategies that are based on the realities of the situation to make any meaningful progress. In particular, I would like to offer four principles here, to apply in building real-world healthcare systems that will provide real value:

Keep it simple. Start off by connecting systems one by one. If you're going to propose standards, keep them low in the technology stack; secure email to the pharmacist, for example, is a better first step than an e-prescription system with all the bells and whistles.ⁱⁱ

“Done” is better than “perfect”. If you don't deliver your system, you might as well not have started. Keep the feature set small at first. Curb your ambitions!

Remember that we're here for patients. Don't insist on features/standards/whatever without clear (and ideally proven!) clinical benefits.

Plan to translate. Standards won't save you from having to write a lot of transformations. Develop a transformation strategy. Get good tools. And hire good clinical analysts.

If you remember one thing from my ramblings here, let it be this: Healthcare data is a lot more complex than data in many other industries. As a result, connecting systems together is also harder and more complex. And, of course: Just because banking did it... doesn't mean it's easy for health care.

Notes

ⁱ I first saw this example, and many others like it, in David Markwell's excellent introduction to SNOMED CT. I highly recommend attending one of his sessions if you are interested in this topic: http://www.hiqa.ie/system/files/workshop_HI_CicSnomedCt_20100127.pdf.

ⁱⁱ For an example of using simple technologies to solve real HIT problems, check out the Direct Project at <http://directproject.org/>. They are using secure email to replace fax machines! An excellent overview can be found here: <http://wiki.directproject.org/file/view/DirectProjectOverview.pdf>.