HIT Standards Committee
Clinical Operations Workgroup
Transcript
August 29, 2013

Presentation

Michelle Consolazio – Federal Advisory Committee Act Program Lead – Office of the National Coordinator
Thank you. Good morning everyone, this is Michelle Consolazio with the Office of the National Coordinator. This is a meeting of the Health IT Standards Clinical Operations Workgroup. This is a public call and there will be time for public comment at the end of the call. The meeting is being transcribed and recorded, so please remember to state your name before speaking. I'll now take roll. Jamie Ferguson.

Jamie Ferguson – Kaiser Permanente, Institute for Health Policy
Present.

Michelle Consolazio – Federal Advisory Committee Act Program Lead – Office of the National Coordinator
John Halamka?

John Halamka, MD, MS – Chief Informatics Officer – Harvard Medical School/Beth Israel Deaconess Medical Center
Present.

Michelle Consolazio – Federal Advisory Committee Act Program Lead – Office of the National Coordinator
Martin Harris? Chris Chute? Donald Bechtel? Liz Johnson? John Klimek? Joyce Sensmeier? Kevin Hutchinson? Cris Ross? Becky Kush? Wes Rishel? Dan Vreeman? Stanley Huff? Marjorie Rallins? Floyd Eisenberg?

Floyd Eisenberg, MD, MPH, FACP – Independent Consultant
Present.

Michelle Consolazio – Federal Advisory Committee Act Program Lead – Office of the National Coordinator
Hi Floyd. Jeremy Delinsky? Kim Nolen?

Kim Nolen, PharmD – Medical Outcomes Specialist – Pfizer, Inc.
Here.

Michelle Consolazio – Federal Advisory Committee Act Program Lead – Office of the National Coordinator
Jay Crowley? Karen Trudel? Nancy Orvis? Terrie Reid? Clem McDonald? Marjorie Greenberg? Kevin Brady? Are there any ONC staff members on the line.

Farrah Darbouze, MPH – Program Analyst – Office of the National Coordinator
Hi, this is Farrah Darbouze from the Office of Science and Technology.

Michelle Consolazio – Federal Advisory Committee Act Program Lead – Office of the National Coordinator
Thanks Farrah. And I'll turn it over to you Jamie.

Jamie Ferguson – Kaiser Permanente, Institute for Health Policy
Great. Thank you. Well, as you may have heard, this is I guess one in a series of continuing meetings on image-sharing, which we've been asked to consider. We're very fortunate to have with us here today David Clunie, to add his own perspectives on a variety of standards related issues having to do with the use of DICOM and other related technologies. And so I'm really looking forward to this discussion to add I think a different or fresh perspective on the potential use of many of the same pieces of the puzzle that we've been discussing in our previous meetings. John, do you want to add anything to that?

John Halamka, MD, MS – Chief Informatics Officer – Harvard Medical School/Beth Israel Deaconess Medical Center
I know David and I have exchanged many emails, I did get some valuable perspective on where the standards are evolving? Are new standards necessary? Are constraints on existent standards sufficient? What are possible architectures? And David, I was at the University of Chicago on Monday and Tuesday and spoke with several colleagues there who have worked on some of these standards and projects and really look forward to your comments. Because as you know, as we've been gathering this testimony, we've had a tension between their installed base and existent standards, the evolution of more Internet web friendly standards, how do you both engineer for the future, but at the same time, leverage the installed base. And are standards the issue or is it, in fact, we just need better guidance, more constraints. So, I hope you answer all these questions.

Jamie Ferguson – Kaiser Permanente, Institute for Health Policy
Yeah. And, so this is Jamie. Just one more thought before I turn it over to David and that is, I'd really like us to consider and think through, as we go through this discussion today, areas or whether there are or which areas there may be in this discussion where rigid regulatory standardization may, in fact, be counterproductive, considering the installed base and the pace of innovation. So, I think with that, I'd like to just turn it over to David and let's get started.

David Clunie, MBBS, FRACR – PixelMed Publishing
Okay. If you could bring up the slides, whoever's in charge, yup, thanks. And those are exactly the topics I had planned to cover, so hopefully I will do justice to them and if not, you can interrogate me at the end. So, I only have about 19 slides or so. By way of background, I'm a radiologist theoretically, neuro-radiologist, but I haven't practiced in a while and now I focus exclusively on the medical informatics aspect of things, particularly DICOM, which I've been doing since the early '90's and currently the editor of the DICOM Standard and also IHE. I was a bit of a latecomer to IHE, but I'm currently the co-chair of the radiology and technical – so, hopefully I can bring some kind of veracity to the discussion about DICOM and IHE and where standards are and are not necessary. And also if you have any direct feedback that we need to incorporate in either standard or set of standards, we can take the feedback from this group back to DICOM and IHE quite easily.

So if we look first at the question of use cases, I've been listening to the testimony up to this point and the discussion and I think the best way to frame this is with the view, download and transmit framework that's been discussed for meaningful use, use that in the context of images and image-sharing. And define what for us each of these terms means. Viewing being selecting the images to be viewed, navigating amongst them, displaying them, potentially interacting with them to window and zoom and so on. Perhaps measuring them, perhaps analyzing them, perhaps that's too – that's beyond the scope those are subjects for discussion. Downloading to the local machine, perhaps to use locally or to record on media for both using them directly, for archiving them and sharing them is obviously an important use case. And transmitting them to a third party, whether it be from provider to provider or to an archive or to some kind of third party analysis service, like a CAM or CAD service, manufacturing or diagnosis detection service. Those are all possibilities that need to be considered.

And for each, if we look at the who, what, when, where and why and how, I think that encompasses most of the things that have been discussed up until this point. And by taking this more general approach perhaps, rather than just looking at the point-to-point clinician to patient, imager to patient, imager to clinician use cases, we can come out with a commonality in terms of mechanisms and standards and requirements. And the who encompasses not just the people who do the imaging or who order the images, or to whom the patient is referred for further treatment, but also as has been pointed out, this concept of a clinical team or patient care team that may span traditional intra-enterprise boundaries needs to be considered and also the empowerment of the patient themselves to both manage their care and take responsibility for some aspects of this. I mean, that's also an important consideration, obviously.

What needs to be shared? The complete set, a subset, only key images, we don't know if the report is included, whether or not other –ologies, cardiology obviously is usually included along with radiology, but things like ophthalmology, dermatology, gastrointestinal endoscopy, all those other potential users of images need to be considered in this context, and not just looking at it from a radiology-centric perspective. When does this occur? Is manual, is it sort of Blue Button approach where the user triggers this or is it automatic in that there are pre-defined triggers that cause these actions to transpire? Both mechanisms obviously need to be supported and the list of triggers need to be enumerated.

Where do the images live and from where is the transaction initiated, from the EHR or the PHR? Do the images live in a PACS or a VNA or some kind of HIE based image archive or some other architectural solution? And why is one sharing? And one has to envisage that there's a spectrum of different purposes for which one is sharing, but whatever is shared needs to support all of the potential use cases since, at the initial time of sharing, one does not necessarily know exactly what subsequent uses are going to be. It may be years down the track before you understand why the images are necessary, particularly where they're being used as priors for comparison.

And of course what we're most focused on in the standards community is how to do it? Whether to use a push or pull mechanism or architecture? Whether the payload needs to be standardized and with what standard? Whether the protocol needs to be standardized and with what standard? And what quality and speed tradeoffs are necessary to solve each particular use case. And also, of course, which identifiers are used to recognize that these images belong to this patient in this domain as opposed to this patient as identified in a different domain. Next slide please.

I think actually, by the way, some of the bullet points on my previous slide were cut off on the PDF that you're seeing on the screen, so, I did share the PowerPoint so perhaps afterwards, that can be distributed to anybody who's missing a couple of the points. On this particular slide, which is entitled AMA Safety Panel, I wanted to bring to your attention a couple of the lessons we learned from the CD standardization exercise. As you probably know, DICOM CDs have been around since the mid-'90s and they first started off in cardiology as a replacement for 35 mm cine film, a role in which they've been very successful. But over the years, people have become disenchanted with them for very specific reasons, and one of the reasons is because people used to write non-standard format CDs and you were dependent upon the burned – the viewer that was burned on the media, which might or might not execute in your local environment. And furthermore, sometimes the content of the CD wasn't complete or wasn't of sufficient quality to get the job done.

So, activity initiated by a group of neurosurgeons and implemented through the AMA ultimately resulted in the Safety Panel, which came to a consensus of both clinician groups, a broad range of clinicians, not just the neurosurgeons, as well as representatives of the imaging industry. And they ended up stating the following, "All medical imaging data distributed should be a complete set of images of diagnostic quality in compliance with those found in the IHE PDI Profile. So in summary then, it had to be complete, they had to be diagnostic and they had to be standard. And it's important to recognize that in this context, the clinicians meant diagnostic in terms of their decision-making. Because ultimately the clinician is making a medical management decision for the patient. Radiologists like to think that they are making the diagnosis, but ultimately clinicians are making the diagnosis, based on input from the radiologist, report is just one source of that. So, one has to be very careful in making a decision to compromise the quality of images that are distributed after the radiology report has been issued because there – as we all know, there's more to it than that. Next slide please.

So a whole bunch of other lessons were learned from the DICOM and IHE PDI CD experience. Over time as various different forms and variants of the CDs proliferated, IHE decided to get into the act and release what they call the Portable Data for Imaging, or PDI Profile, which essentially constrains some of the aspects of how DICOM CDs may be encoded, with the goal of simplifying reading, displaying and importing tasks. And in this context we're talking not just about CDs, but also DVDs and conceivably even USB memory sticks and external hard drives and the like. But, one aspect of this was to address the question of the on-board viewer, and we went through a couple of iterations of this. For a while we wanted to deprecate the presence of a viewer because of security issues that were arising with executable code. I don't know that anybody's ever actually caught a virus, as one might say, from a DICOM CD, but the theoretical possibility existed.

And now we've kind of swung in the opposite direction where thanks to the input of the AMA Safety Panel who participated in these IHE activities, we came up with a standard for basic image review, which defined what viewer features were desirable and sufficient for their use case. Also of interest was a concept that IHE PDI introduced which was the optional so-called web content, which was essentially a requirement – well, not a requirement, an option to add HDML plus JPEG versions of navigators and sub – either all or a subset of images that "faithfully represented the patient's clinical condition." This is a nice idea, it was very controversial at the time and we fought about whether this should be mandatory or optional or whatever. In the end we agreed to make it optional, but in the long run, it turned out not to be terribly widely requested or terribly widely implemented. Which is perhaps a lesson for us as we envisage the possibility of JPEG distribution as a network-based mechanism.

The report was also optional for a number of reasons, not the least of which is that the CDs may be created and sent off with the patient immediately after acquisition, long before the report has actually been generated or finalized or signed. But one can include the optional report. The file format was not constrained – is not constrained, and one has to consider the different use cases for the report on the CD, is it just readable or is it actually importable in some kind of a structured form. Next slide please.

I don't want to dwell on this, but as I said, we do have the IHE basic image review profile, which specifies some standard user interface behavior in terms of what kind of interactions are mandated. Direction of mouse movement, keyboard shortcuts, even to the extent of what icons should be used for the various different tools. To be fair, this has not been very widely taken up by the vendors who are not excited about compromising their own look and feel internal standards for user interfaces and they may have a large install base of viewers that are consistent with their PACS user interfaces, but operate in a different way than the basic image review profile specifies. So, this is a bit of a controversial one, so I won't dwell on it, but just so you're aware it exists. Next slide please.

So the next lesson from the CD exercise was the question of importation. As you've heard from other speakers in previous hearings, there are vast numbers of CDs imported and everybody underestimates how many they will be importing once they enter into this. But it has become, now a days, what I might refer to as the standard of care, using that expression a bit carefully. But typically what's done now is for any registered patient who brings media, it's imported into the PACS or the VNA, either for a time-limited or for long-term archival for future use, as radiologists finally get over the fact that the PACS is not for the radiologist's benefit, it's for the patient and the patient caregiver's benefit. And so people are ultimately importing content that's acquired externally and the goal is eventually to have the same user experience, regardless of whether the exam was locally acquired or remotely acquired because ultimately it makes no difference in terms of patient care.

The format issues are largely solved by the DICOM and PDI profile, but there are some specific issues related to import, not the least of which is reconciliation of identifiers, the medical record number, the patient ID or the patient with an external study is very unlikely to be the same as the one that would be used inside the hospital systems. And so one has to deal with that issue, and that's essentially what this workflow profile is all about and it addresses importation of any kind of DICOM content, not just images, what we refer to as evidence documents, structured reports, radiation dose, that kind of stuff. It does not address the import of non-DICOM reports at the present time. Next slide please.

So to pick up on a few of the standards related issues that are related to the use cases that I enumerated initially. First of all, there's a question of network sharing and what the payload should be. So if the payload constitutes a complete set of DICOM images, then that satisfies the required quality standard of the most demanding use cases, and it also allows for all of the import, read and analysis use cases. And furthermore, it facilitates conversion into something that can be used for other purposes that are perhaps less demanding. So inside the enterprise, where there is transfer from the modality to the archive or server, or even in the standalone clinic where you're acquiring images from a modality and sending them to a remote server, DICOM is obviously the way to go and there's not much debate about that.

For inter-provider transfer again, DICOM is fundamental and necessary because of the need to be able to support any of the use cases that the recipient might need to address. So for either point-to-point transfers, which might be considered as the view, download, transmit "transmit" component, or for transfers to another provider that are mediated by a third party, like the patient carrying the information or initiating the transfer, the "download" use case if you like, both of these really should be addressed in terms of providing the capability to download a complete set of DICOM images. Not that that's necessary in every use case, but the ability needs to exist.

For viewing though, the payload can be any suitable format for the task, depending on how the receiver and provider are linked together, and I'll come back to this and address it in much more detail. But in summary, one might argue that the best payload for viewing for demanding tasks is DICOM and for simpler tasks is perhaps JPEG or PNG or GIF or some other consumer format that is appropriate for the recipient. And I hesitate to equate these to the diagnostic and review use cases. From the 30,000 foot view, these might seem to be a 1:1 relationship, but it's not entirely clear, for example, that ordinary JPEG images are not perhaps diagnostic under some circumstances, or for some review tasks, particularly those that involve measurement, for example, orthopedic templating is a typical one, you may need the additional metadata and pixel data that are present in the DICOM image. Next slide please.

So what about the protocol? And to put it simply, who really cares as long as it works? And I think to Jamie's point, standards are not always needed or appropriate when you have – particularly when you have a tightly coupled client server to use that expression, implementation, which may be appropriate under some circumstances. And different protocols may be required for each of the view, download and transmit scenarios. And the selection depends mostly upon which actors are involved and what their relationship is – and what their geographic relationship is, what the distance is between them, are they inside the enterprise or are they outside the enterprise. I think Hamid last time spoke about the local versus long-distance analogy in this respect. So, for example, with respect to which actors are involved, if the EHR is actually performing the view, download and transmit, then you may need one set of protocols, whereas if it's delegating the activity to the PACS of Vendor Neutral Archive who's performing the work on behalf of an EHR-initiated request, you may need a completely different set of standards. And likewise, different protocols inside the facility versus to a partner versus to a relative stranger. Next page please.

So if we look at each of the different view, download, transmit scenarios, with respect to the protocol, if we think of transmit as being a simple push, then the DICOM – the original DICOM C-STORE that one uses traditionally inside the hospital, actually works fine for pushing beyond the enterprise as well, if the other end is set up to listen appropriately, and that can become a challenge, particularly in terms of addressing. And so to provide a mechanism for doing this from within an HTTP browser-oriented environment, we have a new service in DICOM that's just very recently been finished, which is called STOW, which I think is Storage Over the Web – or I can't remember exactly what the abbreviation stands for, and it's STOW-RS because it's based on RESTful Services. And simplistically put, this is just an HTTP POST of DICOM images. What – the particular use case that this enables is to be able, within a browser, to create a new image, a derived image or an annotated image or an annotated object, and initiate an HTD post from within the browser application. And it also enables the reuse of a whole bunch of HTTP infrastructure that may be between the sender and the recipient.

From an IHE perspective, we have XDR-I which allows you to reuse the XDS push transactions to call them that, without the need for an XDS-I manifest, and I won't go into detail about what that is unless you want me to. And this is in some ways related to XDM and Direct and you probably know plenty about that, which I don't need to go into, but the possibility of email via SMTP also arises, although scaling SMTP to handle many hundreds of megabytes or gigabytes of images doesn't necessarily provide an ideal solution. Regardless, in the push use case, the transmit use case, clearly the sender and receiver do need to agree upon standards. So this is one of those scenarios in which agreement on which standards or which of – what subset of several standards does the receiver in particular need to implement then. If I were a vendor implementing this kind of scenario, I would obviously implement pretty much everything so that I could receive anything from anybody.

One does need to address the question of who initiates this, and as I mentioned before, if this is not done directly by the EHR, how do they delegate this to somebody who performs it on their behalf. And I mentioned the need to solve the addressing problem, which you guys have been over with the Direct approach, in terms of the email addresses and certificates and so on. But, for any of the other mechanisms, you need some kind of URL endpoint or in the DICOM case, you need the IP address port and AE title and so on. Next slide.

The pull use case, the "download" use case of VDT, similarly DICOM has both C-GET and C-MOVE services. And C-GET and C-MOVE work fine within an enterprise, C-GET actually works fine over the enterprise too because there's no requirement to open a hole in the firewall for an inbound data connection. That said, it's not terribly widely used for that so as a consequence, and given the preference for many HTTP oriented services, the first HTTP service to be introduced at DICOM in the early 2000's was the URI based WADO web access to DICOM updates mechanism. And that's since been extended to embrace some of the IHE transactions, which is WADO web services and more recently, the RESTful equivalents, which is WADO-RS. These are all officially part of the DICOM standard now and to put it simplistically, these all allow some form of HTTP GET of either the DICOM original image or the DICOM image rendered as something more consumer friendly, particularly as a JPEG. And it allows certain parameters to be specified to make that happen in the appropriate way under the control of the requester. The more recent versions also allow for separation of the metadata from the pixel data, as has been mentioned by previous people who have testified and various combinations of these services allow either single or multiple images to be retrieved.

IHE XTS-I, you've probably heard of, I'm sure you're familiar with XDS which is a document oriented transport architecture, if you like, for having both a registry of information, a repository where the information lives. In the imaging case, initially the thought was we could just stuff all the images into the repository, just like any other document, but the repositories weren't really designed to scale to thousands and thousands of images per study, for example. And so an extra level of indirection was added below XDS, which stores the manifest as a document in the repository and then there's an extra step to retrieve the images identified in the manifest from the imaging document source.

Finally, there is an entirely reasonable approach, which is to have a very proprietary, tightly coupled client server implementation where, particularly if you're operating from within a web browser or a thick client. As long as the web browser code or the thick client are provided by the server, then it doesn't really matter from your perspective, accessing it as a user within the browser, how the images get downloaded is less important than the ability to save them to a file, if you like. In that vein, one also needs to consider that in the specifications – the functional requirements for the download part of view, download and transmit. One definitely wants to be able to download as the original DICOM images, but it's also of great utility to be able to download as some other format, particularly selected images, and to download individual images perhaps, key images as JPEG or whatever. Next slide.

Viewing, on the other hand, is a much more complicated use case, transmit and download are easy by comparison. And there may or may not be a place for standards in this particular setting. It depends entirely on the viewer technology and the paradigm, or in other words, who provides the viewer code, particularly if it's being provided by the web server that the standard browser is talking to. And you'll have heard the term zero footprint, which essentially means no help or applications, no plug-ins, no applets, no Flash or Silverlight and maybe not even any JavaScript. And I use the term absolute zero to refer to the most minimal set of web technology that you need to enable this to happen, which might include HTML before HTML 5, user frames and tables and images that the browser is capable of rendering directly from an image reference. Whether or not one includes JavaScript in this set of capabilities it's – I vacillate about that personally. There are still people who turn off JavaScript for security reasons, so, you may want to have a viewer that's capable of functioning without even resorting to JavaScript.

And then there's almost zero, which might include JavaScript, might include HTML5 Canvass, although that's not universally available yet and it's remarkable how many enterprises still have older browsers and operating systems they need to support. And then there are things that are sort of pretending to be zero that have a Flash or a Silverlight dependency. I don't mean to pick on any particular vendors technology, but we've seen kind of a trend towards Flash and then a trend away from it as mobile devices, which don't support Flash, have become prevalent. And then you have things that are not zero at all, which are just fine for many deployments, may use Active X controls or depend on a particular viewer – correction, browser, or I use a Java applet or whatever. You may even have a thick client that's being spawned by the browser and the mobile device applications and EHR applications are not always just passive pieces of code that run within the browser on the mobile device, they may well be handwritten code that is optimized for the scenario. And it's also important to bear in mind that people have been doing web-based PACS and remote viewers since the mid to late '90s and they've become de rigor in terms of a PACS purchase. So it's not like this is inherently new technology, but it's certainly involving a new direction, particularly as device capability and mobile devices become prevalent, and I'll talk about that a bit later. Next slide please.

So if we look at what's entailed in a tightly coupled client server or browser server, if you want to call it that, interface, where do standards fit in or not? Typically these are web-based including, but not limited to variance of the zero footprint scenario. The server has the images or is a proxy that is able to obtain them and serve them up in the form that's necessary. Arguably, no standard protocol for communication between the browser and the server is necessary, for example, JavaScript can HTTP GET anything, so what that anything is and how that interaction transpires is somewhat academic. And if you're using a code beyond the browser, like a plug-in or an applet, then you can use any kind of transport mechanism you like. We often refer to this as server-side rendering, which can even enable 3-D or advanced visualization tasks, whether they be volume rendering or surface rendering or CT-PET fusion or whatever, all of this can be done interactively with the server. And arguably no standard payload is needed either in that JavaScript can pretty much pose as anything these days, and create the necessary images for rendering on the screen, as you want on the fly. And this even extends to including DICOM – and interpretation in JavaScript if you want to go that direction.

Typically consumer oriented formats that the browser can handle directly from an image reference may be used, especially if there's no interactivity required that, for example, requires windowing or zooming or panning with high quality interpolation. On the other hand, if the viewer server is decoupled from the image source, there is a strong argument for using standards because then the code downloaded from the viewer can access images from another source – from another server. And there it makes sense to choose a standard HTTP-based protocol and this is specifically what the WADO-URI protocol was designed for, both to allow viewers to reference images that exist on a different server as well as to follow links to URIs that were perhaps embedded in a report, for example. And in this vein, there is discussion of the so-called universal or clinical viewers, which have become new buzzwords in the medical imaging industry of late. And one aspect of these might be that they are image source independent and use standard protocols, whether it be WADO or XDS to access those images, even though the images may reside on different archives from different vendors and be rendered in the same viewer, which is still delivered via a web-based mechanism from a web server, independent of the images. Next slide please.

And then finally, the other dimension of view as a challenge from a standards perspective is the question of separation of the requester from the performance, in other words, the ability to delegate the ability to view without having the mechanism to perform a viewing oneself. So you can imagine that an EHR user might want to view a study, and the PACS might want to actually perform it by providing its viewer. We need to find a way to communicate that request to the PACS and make it do it. EHR vendors have made it very clear they do not want to store all of the images, and that's fine. And there would be little point in them performing such a redundant function and it takes some skill to archive images properly, and a very common proprietary pattern is to have customized interfaces between every EHR and every PACS that that EHR might talk to, to address this problem.

One mechanism, for example, is encrypted URLs which both identify the images, authorize the EHR user to do it and this authorization may be, in some respects, time limited, and that's communicated in a proprietary manner through a URL. Unfortunately this means that there's "n" times "m" permutations of EHR and PACS, each of which has to be customized to address this. And the number of PACS vendors is measured in the dozens globally perhaps but the number of EHR vendors is measured in the hundreds. And when you include PHRs as well as EHRs and all other permutations and combinations of folks who might want access to the images, this doesn't scale very well, a good source of revenue for EHR vendors customization perhaps, but not very scalable. Another option is to store fully qualified links, URLs, directly to a particular viewer implementation, but the problem with this is that these go stale over time. And one can imagine that as technology refreshes and organizations come and go, merge, fail and so on and so forth, static links are going to be a problem.

Common identifiers, dates and so on, when the study was performed or around which the study was performed may be more reliable. So to this end, we recently finished a new profile in IHE called the invoke image display profile, which is basically just a standard display request. And so that theoretically, if everyone were to adopt this, would reduce the number of permutations to "n" plus "m" as opposed to "n" times "m." Next slide.

So I won't go into detail, but the invoke image display profile is a minimalist means of image enabling non-image aware systems. And it uses the simplest possible HTTP-based request, everything's stuffed into a URL and a couple of parameters. It supports both patient and study level indication, you can say view all studies for patient or just the most recent studies for a patient or this particular study, but it is usable without a priori knowledge of individual study identifiers, but doesn't preclude you identifying a study by GUID or its accession number or whatever. It does require the server fulfilling the display request, at the request of the user, to perform what we call interactive viewing, which includes scrolling, navigating, zooming, panning, windowing and so on.

It does require the server to be capable of delivering either review or diagnostic quality images at the request of the user. What that means is somewhat open to debate, but diagnostic quality is a requirement, and it does provide the ability for the user to request the display of only the key images, those that have been designated by somebody, typically the interpreting radiologist, as key to interpretation of the case. And this is completely independent of how and where from the server gets the images or where they're stored.

So in other words, this could be provided as a wrapper to an existing entire enterprise of PACS or it could be built directly into the PACS or VNA or some other image archival system, centralized or localized. And it can use any mutually agreed upon HTTP security mechanism. If there's one thing we know in IHE, it's that coming to consensus on security mechanisms is hard, and I'll come back to that briefly. And so this mechanism is designed to be independent of that – of the choice of mechanism, but dependent upon the choice of some appropriate mechanism for the use case. Next slide please.

I've alluded to mobile devices several times and as we all know, these have relatively limited memory, CPU and network bandwidth considerations to take into account. On the other hand, these rapidly increase over time, so one can do more and more every week practically. One of the important aspects of medical devices is not to assume that they are only suitable for low quality use cases or use cases where one can only use them in desperation, in order to make a quick decision because one doesn't happen to have access to a "real system." And there are FDA cleared mobile applications for diagnostic uses and this may be the trend.

When you get into the mobile world though, as I think Jamie was alluding to at the beginning of this session, the technology advances so rapidly that one needs to be very wary about getting too dogmatic about which standards to use. One enters into the RESTful versus SOAP religious wars, the JSON versus XML metadata religious wars, which are rooted in performance requirements as well as simplicity requirements and tooling support requirements available on the platform. Not all browsers on mobile devices support HTML-5 and Canvas elements, so you need to take that factor into account.

And there's a whole new crop of MHD, mobile health device standards mirroring XDS capabilities that IHE has come out with. These are new, it remains to be seen what limitations they have with respect to the use cases that XDS can be used for. But, I would predict that very soon we're likely to see a whole family of mobile device friendly RESTful JSON oriented MHD standards that provide a bridge to an existing XDS infrastructure or completely replace an existing XDS infrastructure by using the same functional specifications.

In terms of the payload or what's transferred, one can span the gamut from full DICOM images through JPEGs through proprietary formats that are optimized for the particular transport task required. Same for the protocol, you can do DICOM off a mobile device or you can do WADO off a mobile device or you may choose to use a proprietary, typically HTTP GET based interaction. One of the key aspects is the control of the viewing environment and the display quality and the FDA had several concerns about this when these applications first started coming to light. And those problems are being addressed and there is kind of a paradigm for handling that now that the FDA and the manufacturers have established. And one day it may be that all viewing is performed on remote – on mobile devices that are wirelessly connected, we shall see I suppose.

And just to address a comment that John made, I think at the last testimony by Hamid and Mike from lifeIMAGE, there are indeed mobile applications that do execute a traditional DICOM interpretation, use DICOM TCP/IP transactions to provide a very capable viewer that runs on a mobile device. And there are other applications that do this service, do the – rendering and provide the information to the mobile device, inside the browser on a – or a thick application on the mobile device. And I have pictured examples of each here, and I won't go into naming any vendors here. Next page please.

The question of architecture has been raised and traditionally we group these into pull and push and in centralized versus distributed and various permutations and combinations of those. In a push architecture, particularly in a point-to-point setting, it's kind of easy and tempting, it may result in duplication, the same object stored in many places. But I think as Dave Mendelson mentioned on a previous occasion, that's not necessarily a bad thing, but it's no substitute for a formal archival strategy, of course. And it also raises the question of change management issues, what if the images are incorrect, they were for the wrong patient, the side was wrong, that kind of stuff, and they're corrected locally, how do you make sure that they're pushed out and all corrections are handled downstream everywhere where the images may have gone already.

Pull architectures are divided into various different categories, depending on whether the discovery part of it, the queries are distributed or federated in some manner, or whether there's a centralized registry. Similarly, obviously the images can be stored in some kind of central place or exposed locally at the edges of the sources of the images. In the last case, the local storage case, links do go stale, enterprises do go out of business and the question is, where are the images going to be then, even though the patient may still need to be cared for. And there are various hybrid combinations where one has brokered or clearinghouse style transfers where some kind of intermediary holds the images transiently, possibly in an encrypted form. The sender pushes them, the recipient is notified and then pulls them, in some manner, for example, following the URL and the email or whatever. And we've all used things like Dropbox and so on for sharing files, which operate in this kind of manner and you've probably seen Filelink and its ilk for handling large email attachments rather than sending them via SMTP. Next page please.

Some of the other considerations I just wanted to go over briefly are, if we have time, obviously there are business model and sustainability issues that are directly tied to the choice of architecture, and the bottom line being, who pays. And for some architectures this may be insurmountable under some circumstances, and that's not a matter of standards, but we need to take this into account when we start to prescribe which standards will be used or which architecture will be used or we need to get – avoid being too prescriptive in that respect. And we can learn something from the global experience. In Canada, as you probably know, in the various different provinces they have regional centralized repositories where the images live, as opposed to depending on the local archives being exposed. In the UK, the image exchange protocol, IEP, arose as a consequence of failure of the national effort to satisfy the trusts. And initially a point-to-point push mechanism was used, which was quite successful, and then they became brokered through the portal and now there is a trend towards expanding IP to become a centralized registry if not repository of images.

The question of whether the report is in scope or not is influenced by a range of different factors not the least of which is the lack of standardization of a format for the report. There are too many standards and not enough incentives for people to offer structure. Ultimately though, the report is just another document – identifying it. And the question arises, should the report be dispatched and distributed and handled separately or should we leverage the convenience of packaging it with the images, like we sometimes do on CDs. And if we do provide redundant pathways for transport of a report, how do we handle questions of duplication and amendments and inconsistent amendments. Amendments to reports are made perhaps a little bit more often than images, images are not so often changed. Although when they are, when it's the wrong patient or the wrong side, it is a critical issue, as I mentioned. Next page please.

So some of the other considerations, and one of the most vital ones of course, is everything we lump under the heading of security, be it authentication, authorization, access controls, single sign on, trust, all that kind of stuff. None of this is really image specific and ideally we would just leverage whatever mechanism the EHR community finally standardizes upon in this respect, as we know, that's been hard. The additional twist with imaging is that if the EHR is going to request a separate system to provide images, either for download or transmit or for viewing, likewise for patient portals, the EHR portal may be separate from the PACS portal. Additional security questions come in in terms of delegation and trust between the EHR and the image aware system, and that's perhaps an extra requirement on the EHR security mechanisms.

Identifiers, I think several people who have testified previously have talked about how nice it would be if there were a single national provider identifier – correction, a national patient identifier. But regardless, in the reality of what we have today, scaling beyond the single site or enterprise is challenging. There is a need to reconcile identifiers and match and map them, not just patient identifiers, but also identifiers of service episodes, accession numbers and so on. That the scalability across enterprises is similar as for any other record, images are not really unique in this respect and qualifying all of the identifiers that are transmitted by the issuer of that identifier, assuming we have a global naming scheme for issuers, would be highly desirable.

And, as you probably know, there's a range of IHE profiles that I won't go into here, but just to enumerate some of them, there's the cross-community access, a variant of XDS. And then there's an imaging variant of that, which is cross-community access imaging, is MINA, which is Multiple Image Manager Archives, which is sort of DICOM – how two PACS communicate and share identifiers that have specific issuers with them essentially. And then there's a range of services PIX, PDQ and PAM which are MPI-related transactions and profiles to handle patient cross-referencing and then graphics and so on, in both push and pull environments.

Glossy image compression, we definitely don't have time to go into today and since Clem is not on the call, we could probably discuss that separately offline. But that does come into question and whether or not it's performed before, after or during any particular phase of the patient's care needs to be considered very carefully. People have a gut feeling that it's okay to compress images and not lose very much, but proving it scientifically is not as easy as it sounds. And of late we've come up with the term diagnostically acceptable irreversible compression, which is a bit of a mouthful, but if you do a Google search on that, you'll come up with some recent papers and reports and standards that pertain to that particular subject and some national recommendations in other countries related to this particular issue.

And then there are other fringe issues that are important to address. The use of standard codes for procedures for anatomy and so on and so forth, this is something that's very difficult to get sites that are internally focused to deal with, and I'm sure you're familiar, very familiar with that issue. And the other one is the question of new features. I've talked about a number of brand new standards, the ink on which is practically still wet, and getting that knowledge of these new standards which were developed for these purposes, out into the field is challenging, particularly for smaller vendors who are not actively participating in the standardization efforts. And so people may regard XDS as the latest thing since sliced bread, but it's been around for probably a decade or more, and yet things like STOW and WADO-RS and so on, are brand new. Same with IID and educating the vendors and the customers and getting them to cooperate and test this, is challenging.

So in conclusion, I would say – next slide please – that we probably don't need any entirely new standards, we've just developed a slew of improvements to existing standards. There – we could certainly do with improved use of existing standards, even the older standards like XDS-I and there is a temptation to attempt to converge on a useful subset of standards. Although as again, I noted that Jamie mentioned at the beginning of the call, that we need to retain the agility to adapt to rapidly changing technology and particularly mobile technology makes this quite apparent. And also despite Hamid's discussion of the local versus long distance experience or scenarios, I think we really need to focus on having a more seamless transition from the local to remote experience. Because ultimately it shouldn't matter where the user is, whether they're in the enterprise or next store at the competing hospital or the VA across the road, or whatever the deal is or whether they're skiing or sitting on the beach, they should be able to perform the same capability because ultimately underneath they're just using the same technology. And the only practical constraint is bandwidth and portability.

Finally, for some functional requirements, proprietary solutions are okay, particularly where there's no interoperability boundary. I mean, the purpose of informatics standards is to cross interoperability boundaries between two different devices that may come from two different implementers or vendors. And if there is no such thing, if the code to perform the task is being downloaded into the browser, and the browser itself is sufficiently standardized, then a proprietary solution is perhaps just fine, as long as we specify what that solution – what problem that solution has to solve and how it needs to – what features it needs to solve it. And finally, keeping it as simple as possible would be nice, if perhaps naïve, but leveraging the huge install base of web-based, particularly web-based PACS and VNA viewers would be highly desirable, rather than reinventing everything from scratch. And I'll stop there.

Jamie Ferguson – Kaiser Permanente, Institute for Health Policy
Wow David, thank you very much. Of course really appreciate the time and great effort that you've gone to pull this presentation together for us. John I wonder if I could turn over the first line of discussion to you?

John Halamka, MD, MS – Chief Informatics Officer – Harvard Medical School/Beth Israel Deaconess Medical Center
Certainly. Thanks so much. So when I was at the University of Chicago Monday and Tuesday David, I was told, you know, why is it that the banking system of America has functioned so well and moved 21 billion transactions, oh, it's because of the ACH network and NACHA as an enabling organization. So the standards, as you point out, are only one piece of the puzzle. So I wonder, as we think about image exchange in the United States and the increasing importance of organization to organizational data sharing, are there enablers that we might need? I mean, so should we recommend to ONC that there is a federal catalyst, whether that's a thing, a network or whether it's an organization, an oversight that might take the standards as you've described them, and ensure that they are more widely adopted for more transactional flow.

David Clunie, MBBS, FRACR – PixelMed Publishing
Well I think the lesson to be learned from the banking industry is that they had an incentive that there was a strong demand for people to be able to use their ATM card to withdraw money from one bank and do it at another's location. And it's not all that long ago, and I'm sure we're all old enough to remember the days when you went to the neighboring state, you had to arrange for a bank draft to even make money cross state boundaries. So to me it's not a question of technology, it's a question of incentives. And the incentives for imaging sharing just aren't there. I mean the state of the art is that when you go to a radiology facility, a freestanding radiology facility or an academic institution and you get your scan, you get your scan and then they might or might not be generous enough to give you a CD at the end of it.

Once upon a time, they'd give you a packet of films, and that was included in the price of the technical fee. Nowadays they might charge you $25.00 for the privilege of getting a CD, particularly if you come back for it for a second time, and there really is no incentive for the facility to do better than that. There's no pressure upon them. Likewise, there's – Hamid and Mike can better talk from the perspective of lifeIMAGE on the business model of sharing via the network and they said they could reduce this down to about a dollar per study or whatever. But still the question arises, who's going to pay that money and I don't think the incentives are really there.

There's no CPT code, as far as I know, for sharing images or burning a CD and until there is, there isn't going to be large-scale motivation to do this. And those "networks" that have been established, have really been established to satisfy specific functional needs and driven by particular stakeholders who have a particular problem to solve, neurosurgery referral networks for example was a case in point, where there is a clear use case and enough motivated people to make it happen. How you do that at a national level and how you do that without funding a central archive, like the Canadians have done, I just don't have the answers. But I don't think specifying standards is – technical standards is necessarily the way to do it. The technical standards will follow the demand and the right standard will be chosen when the right use case arises.

John Halamka, MD, MS – Chief Informatics Officer – Harvard Medical School/Beth Israel Deaconess Medical Center
Very good. Jamie, unfortunately I have to jump to an ONC call with the Chairs, and so let me turn it back to you for the closure here.

Jamie Ferguson – Kaiser Permanente, Institute for Health Policy
Okay. Well, thanks very much. And I do think that we're going to have to have some additional follow up discussion on the points that have been raised in this presentation, not the least of which actually goes back to the beginning of the presentation and the discussion about the payload and the patient safety considerations of the payload. Because I think that the – that AMA recommendation from their Safety Panel would appear to be at odds with some of the previous testimony and discussion that we've heard about the adequacy of using either compression or essentially non-DICOM image subsets. So I do think that we're going to have to come back to that discussion as well.

David Clunie, MBBS, FRACR – PixelMed Publishing
Oh, diagnostic is not synonymous with prohibiting compression, one just needs to be careful about that, so, yeah, we'll definitely have to discuss that more. But I don't want you to leave with the impression that compression equals non-diagnostic, that's not necessarily true if you can prove it's not true.

Jamie Ferguson – Kaiser Permanente, Institute for Health Policy
No, I understand that and you mentioned a number of papers and the work that's been done on the diagnostically acceptable irreversible compression, so certainly, we just need to under – really understand that, I think, better.

David Clunie, MBBS, FRACR - PixelMed Publishing – DICOM, IHE, PACS and Imaging Clinical Trials Publishing, Consulting and Software Development
Sure.

Jamie Ferguson – Kaiser Permanente, Institute for Health Policy
But unfortunately we are out of time for this meeting and so I think we have to see if there are any public comments.

Public Comment

Ashley Griffin – Management Assistant – Altarum Institute

If you are on the phone and would like to make a public comment, please press *1 at this time. If you are listening via your computer speakers, you may dial 1-877-705-2976 and press *1 to be placed in the comment queue. We have no comment at this time.