Thanks Don. I was particularly pleased to accept the invitation to present at this opening talk because back in 1999 of the APM newsletter, there was the statement that the absence of medical physicists in the PACS area was worrisome. PACS is fast becoming the core technology around which all radiology departments will be organized and that was in an article that the current article, Don Frye and I wrote back in 1999. We were concerned at that point because while both had pretty much automated practices or electronic practices and the site visits, from two different vendors, the site visit teams hardly ever had a physicist involved and so we think this is an important area because we felt that even back in '99 that the medical physicists, particularly the clinical medical physicists, is uniquely qualified by training and experience to design, implement and direct the electronic radiology practice at least at the beginning. And then after that to be

able to serve as a consultant to help the institution through some of their growing pains. PACS is but a portion of the whole electronic radiology practice, but it's intimately involved with the electronic medical record or the hospital information system and the radiology information system. Generally, PACS only concerned with or folks think of it only concerned with image acquisition or the management of interpretation. We think one of the important features here is that it is not a stand alone part of the enterprise. In fact, once PACS is enable in an institution, images are used far more in that institution than ever before largely because they're easy to get to. We could've retitled this actually as the ongoing relationship of PACS with the HIS, the RIS and other TLA's and of course begs the question what's a TLA. A TLA is a three letter acronym for three letter acronyms. Of which there are lots of them in PACS and would encourage those of

you to whom this technology is new that you hear an acronym not defined, please ask that person to define that. Well as in many other fields, if there's a TLA there's probably a degenerate case and that's true in PACS as well. There are FLA's, which are either four or five letter acronyms. And there are lots of those as well. Not uncommonly new fields develop a new lexicon. We're gonna talk a little bit about the motivation of why one would get into PACS, some overview of architecture issues and interface issues. Also what standards exist now and are utilized in PACS and then illustrate how PACS functions by the way which workflow goes in the department. So the first place to start is why would you get involved in this. Originally back in the 80s when PACS actually got started, it was in academic centers because academic centers being academic wanted to do new and interesting things. But that certainly is not the case anymore. We use our

practice at Mayo Clinic Jacksonville as an example and our motivation was quite different. There, our old practice, the conventional screen and film practice, was pretty efficient from the referring physician's perspective. We turned around chest x-rays in about 45 minutes and other exams in one or two hours. That was from the time the patient showed up to the time that the jacket with the report was placed in the referring physician's office. However, that didn't come without a cost. And the cost was that on average about eight people would touch that film as it when through the process. So we knew we had to change. We knew we had to be efficient, we knew that we were gonna have to change some procedures. Change is an interesting thing of course. The picture on the left is radiology at Mayo Clinic, Rochester in the 50s. On the right is radiology in the early 80s or in late 80s rather in Mayo Clinic, Jacksonville. You can see how

much it's changed there, aside from the people; virtually nothing has changed with the exception

of maybe the keyboard. So we knew that there was going to be big change in the department that was for sure. We did replace of course the view box and film with monitors and then we also changed personnel, we have all women radiologists now, no more men. No that's not true, we didn't do it. But we are using speech recognition as you see here now throughout the entire department, we don't have transcribers for anything but the chest and bone area. So the motivation here is to actually improve the quality of the practice while improving the efficiency through expense reduction. Now it's important to understand as I said that PACS comes about not in a vacuum. There are other things going on in the institution. And in our case, there was a desire in the early 90s to automate the clinical practice and that is to use, to get rid of the paper medical record. And it was called the automated clinical practice and we dovetailed and called it the radiology piece or the PACS piece, the automated radiology practice. To give you just an

idea about this one and why we did it, the major reason is to improve efficiency and decrease cost. In the clinic we have a few different types of exams. A comprehensive history and physical, a specialty consultation and a return visit. Prior to automation of the medical practice, these are the steps involved in a comprehensive history and physical. With automation, we eliminated the ones in red. Similarly, for a specialty consultation, these were the steps involved and these were the ones eliminated. Roughly half of the steps were eliminated by automation that is very important to understand, therefore, that if you are going to implement or somebody's going to implement PACS, it must be able to be connected with, in our case, the automated clinical practice or the electronic medical record as well as the radiology information systems. These do not stand alone and we'll see why in just a second. This was the original diagram, this

will kind of diagram if you went to a PACS course. You'd see so many of these that you would get sick of this after awhile because no one can take a look at all those individual elements, but don't be disturbed, these will grow quite quickly. It starts filling up real, real fast. And that's true for virtually anyone's implementation today. And by the way, these kind of diagrams, well in fact, this diagram actually we didn't do all on our own when we started and we started this in 1994. This was done in concert with a vendor and nowadays, by the way, the vendor will help a lot in terms of analyzing the workflow and making decisions about giving you advice on how to configure the architecture itself. Most places do use nowadays, this type of topology though, that is, they are distributed systems and the archive is one piece of that, but it's not very common anymore to find a piece where that is the only node where images are available. The other thing

that I think is real, real important in this endeavor, particularly at the beginning, is that this is a team effort. The whole practice of radiology is not run by just the radiologists or just the technologists or just the engineers or the desk people however, all of these people are important in understanding who interacts with images and the information that concerns the images, all the way down to the person at the front desk. It's a very interesting thing that in most places, the most entry level position is the first person that the patient actually sees when they come in and makes an opinion about what that place is about. So these kinds of things, in terms of automation, we want to think about ways in which we can enhance the efficiency and the job satisfaction of other folks, as well as the radiologists. Now these systems are not all purchased from one vendor. Every vendor that builds these, by the way, will tell you that it will be far, far

better if you buy your radiology information system, your lab system, your EMR and all that from us because then they will all interface together very nicely and you won't have problems. Well, that's not true. It's not true in any of the big companies that we deal with. Our EMR is by____, our PACS is by Siemens. I'm aware of the other big vendors as well. But even internally, the area that writes let's say the RIS software in _____is not the same group that writes the PACS stuff. So in fact, they have interfaces internally. We had, at one time, as many as 28 different interfaces into other systems in order to run the PACS and we'll see why in just a second. It is important to understand though, here that when we're talking about the interfaces we had a rule that we wanted the vendors to adhere to a standard. Largely, either DICOM or HL7, I'll talk about that in just a second or if they could come up with another one that was as

widely accepted. That becomes an issue not on day one, not even maybe three years into it. But it will become an issue at some point in time in terms of maintenance and that's where the cost of the PACS systems becomes large. Now you may not heard of HL7 so I just wanted to sketch out a little bit about what HL7 is. It is a means by which electronic medical records software communicates with other things either internally or externally. It's a fairly large group that got together when this was first established, you can see in addition to ACR, IEEE, HIMS, HIMS is a professional group of IT professionals. There's a good reference for this, it'll be in the handout or it is in the handout, a good reference for this is the site at Duke. HL7 exists at the application layer of the International Standards Organization open systems interface model. And so this is the same kind of model that you'll see in other publications where it talks about communication

between other devices or other software applications. HL7 messages basically have two things. A triggering event and then a message and it looks something like this. This doesn't mean a whole lot to us, but if we rearranged it, we'd see that this maybe is the patient's name, this could be a medical record number, this could be the number for the accession into the radiology information system and then this could be perhaps the folder in which the image is placed, this could be demographics about it and then this poor guy is evidently is gonna get a PIK line put in sometime soon. Now how does this information travel from one to the other? The most popular methodology and you can see some of this out on the floor of the technical exhibits is a broker. A broker in DICOM speak, DICOM is digital image communications in medicine, in DICOM speak a broker is an HISSCP. It's a Health Information System Storage Class Provider. The

details of that aren't important right now, what is important though, is you have a device over here somehow it has to get information about what study was ordered, what's the reason for the study to be ordered and various details about the patient and we'll talk about that in a second. It does that through an HL7 communication to the broker and the broker supplies the device. New radiology information systems get rid of the broker. The whole idea of that is to get rid of that because it's just another point of, it's a potential point where you could have trouble. So if you're implementing a new system you would want a brokerless radiology information system. Now why would you want, well you wouldn't want an RIS interface, right. RIS interface is an interface alright, but it's not one that any of us would really like. We wouldn't want an RIS interface and therein lies the problem. You need a way for the various systems to communicate

with one another so humans are not re-entering data, that's a real, real problem. The interface

itself is the communication protocol so that you can exchange this information either between systems or between software applications, even if they're on the same system, different apps have to talk to one another as this process goes on. Generally, the information that you want is something about the patient, the acquisition device you want a unique identifier for that exam and this is the RIS accession number. The accession number is the unique number given by the radiology information system that is associated with that order. Why this is important, is that when the digital image is then produced at the modality, you would like that accession number to be deposited in group 8, element 50 of the DICOM header, that's where the accession number is. That then ties that image to that report and the report of course, is the medical legal instrument and so therefore, it needs to be referenced. So this accession number is a very important, unique

number and if you don't have that correspondence in that image header, you're going to have to build tables in order to keep this correspondence intact forever. And so that becomes a real, real problem. Just for a second, the information, the RIS communicates with all kinds of other systems. It may get information from registration at the time that the patient comes in or at the time that an order is placed. It has to communicate with imaging modality. It often communicates with the QC workstation that the technologist uses. It communicates a lot, particularly nowadays, with the imaging workstation itself, as well as the archive, the distribution system, most places nowadays use a web distribution system and the HIS or EMR. So this is a key feature and it's a key piece in the overall functioning of an automated practice. I just wanted to illustrate what these interfaces look like by considering an interface with a computed

radiography device. The CR of course, is the most popular replacement for a conventional screen film radiography. You use the, or the interface is in place for a couple reasons. One is data integrity; you wanna make sure that the data do not get messed up from the time of the order to the time the imaging is done. And then secondly, it's a very important feature in terms of the workflow. This communication and this high fidelity nature is not a problem with film and human handling. If a name wasn't right and somebody puts a sticky over it or something like that and writes on it, this is a big problem with an electronic system though. It's not so much a problem when you acquire the data; it's when you wanna retrieve it. There are certain things that you really have to get absolutely right in order to be able to retrieve the data in the future. Things like the identification number; the name is a very important one. Well, names can get messed up.

Somebody might think they were doing an examination on that person, when really it's maybe it's that guy or maybe it's that guy, we're not sure who that is. Well let's assume that we get the name right. Could be Mike Flynn, Michael Flynn or Michael Flynn in caps. Now even though his family thinks that he might be three people, this is only one person. Unfortunately, the electronic systems today will think these are three different folks. So what if we do what we usually do in physics, we establish a convention. We'll keep the names all in caps and put the last name first. Well, FLYNN, Mike, FLYNN Michael, FLYNN Michael J. once again same idea, three different people. So that's the reason that we want an electronic key to this kind of information. It's not so much that acquisition is a problem, it's that the database search in the future won't find that exam. One of the things that makes PACS a greatest utility to most

departments is the fact that they can look at all the old images. It's true at our place and most

other places that I know of, the radiologists are looking at more old images than they ever did before, because of the fact that they are now available and they may not have been before. If you don't do something on a front end to correct this, you're gonna have to do something before the images go into the archive obviously. Studies at our place, both in Rochester and Jacksonville, and at other places have shown that manual data entry of this kind of data comes with at about 15-20% human error. If you did that and you're doing 200 exams a day, which is not a terribly busy department, you're gonna need between 12 to 24 hours or about a 0.2 FTE just to correct all the database errors. So this becomes a big issue. So these interfaces can be done in one of two ways either what's called a modality work list or a query on demand. And we're gonna take a look at a modality work list. This is the console for a CR, a dedicated chest imager, a Digiscan

2T, a Siemens device. And then you can see the list that's provided. That list appears on that monitor in the following way. Every so often, the computer at the CR device goes to the broker and then the broker goes down to the RIS and says give me all the orders for chest exams and then they populate back to the screen. So therefore, at the time of the exam, the technologist just selects, in this case it's done with a plasma screen, just selects which patient it is and then automatically all the headers are located. Now does that mean you're never gonna make a mistake? No. Probably doesn't happen at your place, but we have patients occasionally that answer to the wrong name and so therefore, you can still have error, human error is still possible. However, you won't have those transcription type errors that we saw before. This kind of interface, a modality work list, is available in just about all commercial equipment that would be

out on the floor these days. To get an idea about the interaction of all these groups, we're not gonna go through this in detail, but all of these different systems must communicate with one another and the issue here in this setting, is this broker is the central feature of that. In the future, we will be getting rid of brokers, but it's not something you can do away with right away. That is one thing that would be worthwhile thinking about if you're advising a place on what to do with regard to PACS. To kind of get just an overall wave of PACS functions, let's just take a look at the way the workflow works, at least at our place. The first step is that an exam is ordered. That's a communication between the electronic medical record and in our case, the RIS. After that, the RIS might wanna go ask to say are there any old images. If there are any old images that are pertinent, then I would build a cue, so the day that patient comes, the old images

are available. Then, also, I could download to the modality, a work list and this is the DICOM work list is downloaded to the modality. The exam is then requested and that order is available at the workstation as well as other places within the department. Just before that exam is available for interpretation, we might wanna fetch from the archive the old images, either done previously by a cue or in some PACS now, you can do it on demand because the retrieval is fast enough. The exam is then performed; the images are transferred to the workstation where they are then interpreted. Information from the RIS now goes down to the physician's workstation so that they can report out and they know what the old reports are. It's also very important for the radiologists to look at old exams and old reports. Old exams by themselves are not enough. The radiologist is a bayesian kind of decision maker. They make use of the old information and

therefore, it's important for the reports to be available in a very quick way. Nowadays, a lot of

places use speech recognition with electronic sign off and the report is circulated back to the RIS and then up to the electronic medical record. The images are archived and then distributed throughout the institution. This kind of process that we just went through can take as little, except for the interpretation process, as little as about 5 minutes to do in a fully functioning, automated way and that's an important feature. Now before we close, what about the challenges. What sort of challenges are there? Well, one of the challenges that we have is the amount of data. As we heard in the President's symposium today and it's certainly true, the amount of information that we're producing is astounding. We currently produce about 16,000 images a day for radiologists. Extrapolating our growth pattern at our place, by the year 2007 that will increase to 80,000 images in one day for radiologists. If that radiologist looked at each one of those images for one second, they would be there for 22.2 hours. Well that's not gonna happen

obviously and so, we have to do something about the amount of data. This is a difficult thing sometimes to explain, particularly to a non-IT administrator. But most or a lot of places that are fairly busy produce about 20 gigabytes per day and that works out to be about 5 terabytes per year. Now those terms are not readily known or those people don't really appreciate what this volume is really like. Let's do this, let's call a byte a second and if that's the case, then a gigabyte turns out to be 30 years and a terabyte is 300 centuries. So, those 5 terabytes are like 1500 centuries. So these are massive amounts of data. We live an age now where largely due to Moor's law, the technology has kept up. The issue now is can the humans keep up though. Not can the technology keep up. If you can see our growth pattern there, we've gone off roughly by an order of magnitude in the last seven years and it's accelerating rapidly now for two reasons.

More images made per patient and more patients coming through, so they are both working the wrong way for us. So workstations of the future are going to require a lot more than what we've seen in the past. This is beginning to happen now. Workstations are beginning to employ biometrics, interfaces to the radiology information system, EMR and speech recognition is coming about very, very quickly for both quality, as well as efficiency reasons. The workstations need to be integrated with both reference databases like the RSNA database as well as the internet. Advanced processing is gonna be necessary as we heard about from Dr. Seagull earlier today. Now there have been some early designs. This was an enhancement designed by a chairman of radiology where they put the speedometer over here to measure how fast the individual was reading. But then it occurred to them that there was a lot of wasted time during a

radiologist's day. So that if they could outfit other parts of the hospital with devices that could make sure that images are being interpreted all the time. Well, what is the effect of all this on a practice? And we talked about earlier that the old practice at our place was pretty efficient. Well the new practice became much more efficient. Chest exams are turned around now in five minutes and about ten minutes for the specialty exams. So we've dramatically increased the efficiency. Also cost comes up and I won't dwell on this except for the fact that in many places now can justify a cost model that shows that PACS actually pays for itself. And that certainly was the case for us; this was the estimated savings over the time period from '95 to '99. This was not done by us; it was done by the Mayo bean counters. So, and they're usually pretty touchy about how they allocate the funds. There also is a very important issue regarding quality. As

you implement electronic systems, the quality in terms of the ability to get old images, improves dramatically. As you implement speech recognition, where the person giving the interpretation can look immediately to see whether that is or is not what they really wanted to say. Right versus left, mass is denser or less dense, those kinds of things. That allows you to decrease the medical errors and in this day in age, quality is becoming a very, very large thing and so this is something that if no one else brings up, you should bring up in terms of the reason why somebody would wanna go to PACS. The fact that you have both the images and reports available not only to the radiologist, but too many other people in the institution is a very big improvement in quality. It means that there's less likely to be errors. Now, will they go to zero? No. But if you can help the rest of the institution reduce their errors than that's a very important

and good thing. So that's kind of an overview of PACS from a high level, you might say and the biggest thing and most important point being that PACS does not exist in a vacuum. It must communicate with other systems. There's also these references that are listed in the handouts that give some both general as well as specific references on some of the standards. And there is no question that this will continue to evolve. There's just no question at all. In fact, the PACS that you implement in 2004 will probably be completely replaced by 2008 at the rate they're going now. Thanks very much.