

Welcome to the hail and the hardy, and the proud and the few. It is interesting that finally the whole concept of quality management is becoming recognized as something that warrants a refresher course. Of course, it warrants a refresher course at 7:30 in the morning. That's okay; we know our place. How many out there do quality assurance of any form? Yeah, I thought so. So you know what it is. Yeah. Our learning objectives for today's talk are to understand the principles of error reduction and the procedures for quality management. Whoa! (Ha! Ha!) There's the first example of quality management. I didn't look at this as I copied it from another slide set. Well, you're going to have to know quality management by the time we're done. The first question with quality management, of course, is, "What's quality?" And here's a... here's a definition, which is

often used in industry except adapted to medicine just replacing customer with patient in the first one. Quality consists of those product features (and the products are what we do), which meet the needs of the patient: we have freedom from deficiencies, and conforms to standards and specifications (that is, we get what we want, and that's ... that we're happy with it). Of course trying to, trying to take these concepts and put them into actual practice can be a bit challenging. If you read Zen and the Art of Motorcycle Maintenance, the protagonist in that book has a nervous breakdown trying to establish what is quality. I hope none of you during this talk suffer a nervous breakdown. And who's our customer? As we're thinking about customers, is it the patient? Is it our physician colleagues? The answer is both. You know, we've got to ... we've got to deal with all these people and

have something that satisfies everybody's needs. The concept of quality management is a catchall term, and it refers to all the activities designed to achieve the desired quality in the treatments. And so anything that you're doing to try and make sure you've got the quality falls under quality management. There are two terms that are used kind of loosely, at least here and in a lot of places. Quality control is those activities that *force* a specific quality on a process, that make sure that you get something of the quality you want. Quality assurance, on the other hand, is those activities that *demonstrate* the level of quality in a process. That the quality control forces the level of quality usually works in a methodology on the bottom of this slide, where you evaluate the operating performance, that is, you make a measurement of what something is. You compare it with what you

want, and then you take some action to make sure you got what you wanted. Now having clarified all of that, if you go to the American Society of Quality Control website, they'll point out that unfortunately, people don't always use these terms consistently, and in the literature, sometimes authors will use them exactly in the opposite sense. So as you're doing reading, don't be surprised if you get confused. The authors were also. To point out the difference here between quality control and quality assurance, here we've got some process. And this process we could consider calculating a dose, for example. The process requires inputs. And here we've got four inputs, all of which go into this process. Quality control would go in parallel to the input. That is, they both go into an AND gate. The input doesn't proceed to the process unless the quality control says, "Its okay." And so

each of the inputs has a quality control that looks at it, says, "This is fine, go ahead."

Once the process is done and something comes out the end of the process, you take a look at what comes out, and you compare it with what your expectations would be and say, "Its okay," you can pass that process output along. One of the reasons that there's confusion in what we do as to whether something we're doing is quality control or quality assurance is that the output of this process, as it goes downstream, is usually the input to another process down here somewhere, in which case the quality assurance on this process becomes the quality control for this process as input into the next process. So there's ... there's this dichotomy. And I'm just ... I thought this was so good, we'd start at the beginning again. There we go. This is an example of an error tree, and we'll talk about these a little bit later. An error tree looks at a given error. In this place we have an

error in calculation, which is what we were talking about as our example last time. [The slide changes here but should not until "But if you put quality control into the mix..."] You look at what could have caused this error in calculation. That error in calculation could have resulted from an error in input data, an error in data entry, an error in the calculational algorithm, like the spelling on that, and an error in prescription. Any one of these ... and there's of course a lot more, but just to simplify the slide, we'll take these. An error in any one of these could propagate into the error in calculation. That's why they go in through an OR gate. Any of these would result in an error in calculation. But if you put quality control into the mix, in order to have an error in quality in input data propagate into an error in calculation, you would also need to have an error in your

quality control on that input. So you would have to have your quality control fail, both of these errors going into an AND gate before that error can pass to the OR gate and into the error in calculation. And if you have quality assurance, as this error would go into this OR gate and come out here, the error wouldn't propagate through this AND gate unless you also had an error in your quality assurance. So conceptually, in some process, as you're trying to prevent this error out here, you have the quality control and the quality assurance entering through AND gates to prevent the propagation of these errors. Now we've got them on the slide. I've got them colored green because AND gates are good. They provide protection. And red for an OR gate, because almost always an OR gate is a danger signal implying that anything that goes wrong on that side comes through on the

other side. Quality control procedures, you notice, occurred very frequently on that last tree. There were four of them. And so quality control procedures often require a lot more resources than quality assurance, of which there was only one. And the quality assurance covers a large number of inputs – oh, the quality control – covers a large number of inputs. And so, you're tendency might be to say for efficiency's sake, instead of doing a lot of quality control, "I'll just do one quality assurance at the end." And if everything works well, that's fine. Except if you catch the failure, if you catch the failure by the quality assurance at the end, what that means is you have to go back and reenter all of the inputs, do the – correct whatever was wrong, find whatever was wrong, because quality assurance just tells you something went wrong. It doesn't tell you what went wrong. Go

back, find the problem input correct it, and go through the whole process again. And so

finding an error with quality assurance requires a lot more time and resources to fix than finding the error and correcting it with the quality control. So there's a balance. There's a balance to be found between the quality assurance and the quality control, and in general, they work best together in some mix. The overall layout of quality management looks something like this. There's a lot of discussion as to how it actually looks, but you start with quality planning. It's where you decide what the quality management program's going to look like, what you're going to do. And the quality management has under it the quality assurance, the quality control, and a quality audit that we'll talk about in just a little bit. They all give feedback into quality improvement, which goes back to re-planning your quality management program, and you have a large cycle that goes through this.

The audit is an important part of your quality management. The audit is how you look at what goes on. And it should be an external audit at some point, with an independent, knowledgeable outsider, that's usually referred to as a maven, who can question your assumptions, who knows enough about ... who's knowledgeable enough to ask you about what you're doing, an independent, so that as you've made assumptions, that person hasn't automatically bought in to them. There's two parts, a process audit that looks at what you do, evaluates your procedures, your policies, and a product audit, that reviews your results. In the case of radiotherapy, it's probably looking at your calibrations, your calculations, going through patient charts and seeing what you did, looking at the

outcome of your work. Generally, the audit, to be effective, shouldn't affect the salary as well. As they say, maybe if you do a real good job, they should give you a bonus. But if you ... if they find problems, it shouldn't be considered a deficiency on your part, you should just be able to correct them. Yeah, tell that one to your administrator. In addition to the external audit, you might have an internal audit. You should do this. You should review what you do, and it should be a standing part of your operation to review everything that you do every now and then, and it gives you a chance to question your assumptions and what you've done. As I get older, it gets easier to look at my procedures and not remember the assumptions that made before, and have to justify them to myself again. You should specify your periodicity. If you just say, "I'll do an audit of myself

when I've got time," you aren't going to do it. You should have access to the administration, and talking about what you find in your audits and getting things correct. And of course any audit, whoever is doing the audit, shouldn't have to find blame for what's going on. You're just trying to correct situations. There is a cost to quality management and the quality you want. This is ... this is mostly a picture from the quality control handbook, and it's used in a lot of the textbooks on quality management. It shows costs up here (ordinate) and quality down here (abscissa), where you've got very low quality here (left), and very high quality here (right). This curve is showing the cost of losses, and for an industry, this would be something like if you had a can company turning out cans on an assembly line. Depending on how much quality control you had on

the cans at the end of the line, as you're sending them to get filled, you may have to

throw out cans that aren't the right size, that aren't sealed correctly, etc. And there's a cost to that because you're expending resources needlessly. As the quality goes down, the cost of this waste goes up, and they show it as an asymptote, this becoming infinite at very very low quality, and, of course, as your quality goes up, the costs of waste go down. If you have extremely tight quality control on those cans, you never have to throw out a can, and so you have very little waste in your procedure. However, this is the cost of the quality management program, and at very low quality management, you aren't putting any resources into that program, so it doesn't cost you very much, but as you get better and better, it's costing you more and more in order to achieve that quality. Industry likes to work right here, and the minimum between the cost of losses and the cost of

quality management. And that's all very nice, so you've got some middle quality here. Health care – and I just drew this in as some possible combination curve for health care. Of course we can't really operate in the quality here. A middling quality - it's just completely unfeasible. So our whole curves move over to the higher quality side, and it could be that our cost of ... the cost of failure may go up even faster than in industry. So there's ... we still have to find some balance in that curve, because we can't spend our resources endlessly. We just don't have those.

How often do you perform all your QA and QC procedures? Quality control by definition happens every time the procedure is performed. You look at the input every time and make sure it's right. Quality assurance, on the other hand, has to happen with some

period such that the worst possible conditions for which the quality assurance screens would not produce harm. So you can space out how often you check the output of your machine based on how much it's going to drift, how bad it's going to drift, and how bad the effect on the patient's going to be in between. Generally we check them every day in the morning. We don't check it before every patient because we know they don't tend to drift that much. But you have to find, you have to find in your program, in your planning, a reason that you would do quality assurance with the frequency that you do, some rationale for that. I hesitate to talk about types of errors with a group like this. You've probably learned in high school that there's systematic errors and random errors. So you aren't going to learn anything new with that. But that having been said, I'm going to go

through this anyway. Systematic errors usually tend to be one mistake tucked into a procedure, and they affect all of your patients. It may be something like you forget a factor in your calibration, or have a factor of 2 wrong in your calibration factor on your ionization chamber, and so you – all your doses are off by a factor of 2. Things like that are almost always only found in a process audit, either an internal audit or an external audit. But things like that also should be rooted out. There's no reason – there's no reason that you shouldn't get rid of systematic errors, and they usually are – are things that you can find. Random errors are things that happen on a per-patient basis. These are the things that we're trying to catch through quality management. They'll never be eliminated, because as soon as you – as soon as you put checks in place to catch whatever

errors are taking place, people will creatively find new errors. We're very good at that. I

will point out that the systematic errors usually result from a random error somewhere in the process: that somebody made a mistake, a random error someplace in the process, such as looking up the wrong value for the calibration on the chamber, and that random error then perpetuates as a systematic error. Just one reason that all the steps that go into the process that would be such as calibration should have some checks on them to look for the random errors to keep them from becoming systematic. In general, the quality control – I'm sorry – I'm getting ahead of myself here. Error reduction has two approaches. Error prevention – trying to keep errors from happening, that's the quality control, because it's on the input, you try and keep errors from entering into the process. And as I said before, it consumes more resources. And error interception, which is the quality assurance. And, as I said before, it's a bit riskier, it's a bit riskier to put all your

baskets (?) in looking for - looking for errors at the end of the process with quality assurance. They always work best together in some way. In addition to error prevention and error interception, we also have error mitigation. Error mitigation is saying, "What happens when all of our other quality management fails, and an error happens, and actually the error perpetuates and gets to the patient?" Error mitigation is controlling and reducing the effects of an error, and whatever it is, should be somewhat automatic, shouldn't require people to do anything because you've already missed all this, you probably aren't aware that it's happening. The most common steps in error mitigation is telling the lawyer something happened, and them telling everybody, "Shut up, and don't talk." That's not a very effective method. But here's how that fits in. We had this tree

before, and now we've got the error mitigation here that goes in parallel with the error in calculation, so you don't end up with an error in treatment. Right now in radiotherapy, in diagnostic, in medical physics as a whole, we're getting very good with the quality control and the quality assurance. We do not have very much going on for error mitigation, and in this talk I don't have much to tell you, because it's really just not there. With all those AND gates in the tree, how is it that errors happen and they get through, and that's because very often – very often all those gates are left open, and we'll talk about that towards the end of this talk when we talk about what errors have happened. The first step in designing the quality management program is to do risk assessment, and the first step in risk assessment usually would be to make some sort of a process tree.

This is showing a process tree for low-dose-rate brachytherapy, and while the concepts I'm going to be telling you today are applicable in any part of medical physics. The examples are pulled from brachytherapy because, after all, that's what I do most, and that's what I can - that's what I can tell you about. The process tree – and don't worry about reading the fine print. It is completely unimportant for you. I'll point out what you need to read. Like big brother, I'll tell you what you read. The process goes down the middle, so this is the procedure leading to a patient treatment, and at the end, we've got a successful treatment. We're very optimistic about that. Each of the branches, or the boughs going into the tree are general procedures – here we've got calibration. So calibration is one of the first steps in the tree, although that could be rearranged too. And

off the bough, we have branches, talking about the different steps in the calibration procedure, for example, calculation of strength, calibration factors, readings when you put things in a well chamber, the placement in the well. You can get off any of these branches, twigs, which go into more detail about the one step. This step is entering data, and it has twigs for the factors, the date and the strength of the sources. You can go onto leaves onto these branches. You can keep going into as much detail as is useful in understanding how the process goes. And from this step you also have things like the applicators, your reconstruction, localization. The treatment termination is the last bough. Each of the steps go into this procedure tree or process tree. And you put this together so that you understand what you're dealing with. Next step is to do a fault mode and effects

analysis – FMEA. This shows a sample of a table that one would make for an FMEA, starting with a step of – and I will add the disclaimer that you often have on videos “This slide has been altered to fit the screen,” you would have to have enough space that you could write in any of these – but here's a step in the procedure. And you say, “What does that step – what function does that step perform?” And many of the steps, many single steps can perform many different functions, and so you would have different functions for that given step. And for each of the functions, you would list: What are the potential failures that can happen for that step? What would be the cause of that failure, and each of the potential... there can be many potential failures for any of the functions for that step, and each of the potential failures could have several causes. And for each of them,

you look at what are the current controls in place to prevent that, and given your current controls, what is the likelihood from one to ten that that would occur. That's your “O.” What is the severity (S), if that does happen, from one to ten, and what is the likelihood from one to ten that you will not detect the – that failure (D). And we'll get to the RPN in the next slide. And once again, one for each of those potential failures, you would look at the effects locally, and that would probably be to whatever is going on right there at the time you are doing the step. Intermediate, which is hard to say right in general what that is, but the end result would be in the patient. So any step could have effects right locally with the people doing it, somebody in between, and also in the patient, and you have to consider all of those. So that little end of the table is where a lot of the business takes

place, where we've got the likelihood of occurrence, severity, and the likelihood that you would not detect the failure. Multiplying these together gives you the risk probability number. So it's just the product. The values for each of the variables were between one and ten, with one being very low danger, and ten being very high. In industry, when they see risk probability numbers less than 125, they really don't think very much of that. That seems pretty fine to them. In medicine, I propose that numbers should probably stay down below 20. You might look at things that are less than 20. When the numbers get pretty high, there you've got real problems. And before I move on, I'll just say that there is – the numbers here give you a good idea of how important whatever it is you're looking at is to control, but you have to watch for patterns of the numbers. The product

isn't everything. And if you have ... if you have numbers such as 1-10-1, even though

that would not be likely to occur, and it's not likely you wouldn't detect it, if the severity is a 10, and so you only have a number of 10 here – a very low number, that is something that you should really look at. (*Comment from audience.*) Yeah, please. (*Question from audience: What is the maximum?.*) A thousand – it would be 10 times 10 times 10. And so you're running from one to a thousand on that. Even in industry, when you get numbers up above 500, they say, "Whoa, this is ... this is amazing!"

In the final step as far as the risk analysis goes, is building a fault tree. And we already looked, we already looked at a fault tree when we looked at those inputs into dose calculation error in the example, but this is part of a fault tree for a high-dose-rate brachytherapy. And once again, you don't need to read ... you don't need to read the fine

print here. But I will point out that all this information is on a handout on the meeting website. And you can download the handout and read – actually read some of the numbers – some of the words on this. This – the fault tree – as any fault tree for a process we'd be dealing with, starts with deviation from an adequate treatment, and that could happen either by having a wrong dose distribution or site. You have the wrong dose distribution or you're treating the wrong site, or possibly the wrong patient. Then you say, "What could cause the wrong dose distribution or wrong site?" and that could be that you had a fractionation failure, wrong applicator was used, the treatment planning failed, or the treatment implementation failed. And for each of these, you say, "What could have caused that?" And they have errors that could have led to that. Then you say what – for

each of these – "What could have caused that?" And you go backwards. And you keep going. On this one, let's see, we ... if we went to ... the next ... the next page comes from a different one of these extended, to a dose calculation error. And this is showing all the things that could have led to a dose calculation error. And for each of those, you keep going backwards. And you work and work until you get to a point where you no longer have control over what could have caused that problem. That's the edge of your universe. That's the last thing that you have control over and you're interested in. That would be called the progenitor error, the error that could have caused a cascade through all of this, and eventually hitting the patient. Looking at this tree ... looking at this tree would make the little hairs on the back of your neck stand up and bristle because you have all of these

inputs going into an OR gate. That means anything here that goes wrong propagates. When I've worked with industrial engineers, and they see this type of an error tree, they just say, "Whoa, you guys ... you guys have your business cut out for you." And this is just showing another page. The high-dose-rate fault tree that we put together is ten pages long, so I'm not going to show them all to you. But it just demonstrates the depth that you go in looking at the different errors. So why are we doing all this work? What are each of these steps doing for us? Well, the fault tree helps suggest the placement of QM, that is, where QM is needed, what you have to cover with your quality management. The process tree helps you place the QM in the procedure, and the FMEA helps you pick out which steps you need to address first. Let me ask another question. I started out asking how

many do quality assurance, and it was essentially everybody. How many of you are completely compliant with TG40? Uh-oh. Well, I'll let ... and I ... my hand doesn't go up either. Almost nobody is. We don't have the resources. And the FMEA helps you because you've got to pick and choose. You've got to pick and choose what quality assurance you're going to be doing. You just don't have the resources to do everything unless your department is just over-funded, or adequately funded, however you look at that. And as you're picking, you should have some reason, some reason for picking what you do. [Slide should change here] Well, this is just another page from the high-dose-rate error tree, this one for a dosimetry error. Once again, you don't have to read it – that's not

the purpose. The purpose is to recognize, we've got all these steps just in one, one branch of the error tree. And if we were doing comprehensive quality control and quality management, every one of these steps would be checked. And this is sort of what the TG40 was about. We're trying to figure out what - what would you do to check everything here. And none of us would have those types of resources. There's just an awful lot of boxes on this tree. There are tools that you have to work with. This is ... this is a list taken from the I.S.M.P., which I keep forgetting what they stand for. It's a medical practice. It's actually an organization that's comprised of a lot of medical organizations, and their job is to try to make recommendations for preventing errors in medicine, and they ... they rank types of tools as far as their likely effectiveness in

preventing errors. And at the top of the list are forcing functions, or constraints, such as interlocks, barriers, computer order entry, with some sort of a feedback mechanism to tell physicians when they're prescribing the wrong dose. There's another level, where you've got bar codes, automatic monitoring, computer verification, computer order entry without the feedback, that just ends up clarifying what somebody is prescribing, so you can read the writing. And as it goes down, they become less and less effective. Interesting, the list down here, this is an appendix (the environmental entries). This isn't in their ... theirs. The last, least effective is education and information, training experience instruction. Why is that ... why is that down there last? Because the assumption that they make ... the assumption that they make is people who are doing the procedures have to be trained,

and they have to be knowledgeable to start with. And so this is a given, this is a given that if you take people, put them on the job with the right training to start with, you're still going to have errors. And what do you do, what do you do to fix those errors. And what they point out is one thing that doesn't work is sending people for retraining. Because they already know, they already know what they are supposed to know. They made a mistake. Retraining just doesn't help. It's not effective. And this has been shown in a lot of studies. Also, not all that effective, but often what we can do is making rules and policies. They are nowhere near as effective as an actual interlock that keeps you from doing something you don't want to do. Our tools that we frequently run into, which are fairly effective and relatively inexpensive as far as medical physics practice, for error

prevention, common ... common tools are protocols and forms. A protocol clarifies your expectations and the lines of communication, so that you know what you're doing, what you're supposed to do, everybody knows what everybody else is supposed to do, and

what information goes from one person to the next. Forms ... forms definitely help the process because they format the process. They help you follow what you're supposed to do. They help prevent omissions, and they assist in the transfer of information if you pass the information on a form to somebody. Another common way of preventing errors is monitors, that is, an additional person watching as you do your work. And watching for errors and correcting dynamically as they happen – and I'll point out – it doesn't work. That's not a tool that's effective. It was very commonly the case, when you go through

errors that have happened in radiotherapy and in medicine in general, that somebody has had the job of watching somebody do the work to make sure it's right, and they don't see the errors happen. That's because just watching isn't enough involvement to keep your brain active. And you tend to drift off. Forms should have blanks for everything that's appropriate. If you see that people are making marginal notes, the forms need to be redone with a place for the information people are writing down because if they're writing marginal notes, that information seems to be important to them. And all the information that people do put on the form should be used, or as least potentially used, so you don't make people waste their effort filling in parts of forms that never do anything. Error interception tools, as opposed to the error prevention tools, which we were just

talking about, actually are very similar. Standards and expectations which allow you to judge whether something is correct or not. Forms, again to help prevent omissions and give another opportunity to catch errors, and independent reviews. Now this is different from monitors, because an independent review is having somebody else start fresh and actually look at what's going on and have to do something. We'll look at the criteria on that in just a second. For the standards and expectations, examples would be your previous experience, or somebody else's, assuming that you've verified that they come from correct situations. Existing systems, such as a Manchester system, anything that's been published, they can serve as expectations or standards in evaluating what you've done. The forms, once again, should have a blank for each of the values that you have

that is important for the patient. But they should also have another blank, that is everything should have two blanks – one blank where you put in the value for this patient, and one blank where you put in the expected value. The reason you do that is if you just have a check box that says the value for the patient equals the value expected, it's too easy to just look at them and actually believe that they are the same even though they aren't, whereas if you have to write them down, your brain is more engaged. And as you write the two numbers, you become more conscious of the fact they are not the same. So the forcing of the writing gives you another chance to trigger the recognition of errors. And there's been psychological studies that have shown this is more important for males than females. For some reason, men have to have the feedback of this information going

through your hand. Just like taking notes is valuable even if you never look at them because it forces your brain to follow what your fingers are doing. Women seem to be able to much better just look at things and read them and understand what's going on. Men have a little density problem, I guess. The expected values, of course, should be readily accessible. You don't want to have people having to run around and look up

values. If you have forms on the computer, which is also – it's a very effective way to catch things, but you need to have some sorts of alarms if there's a mismatch between the values that you're entering and what should be there, because otherwise, people get used to just putting things into the computer and letting the computer do stuff and their brain start turning off again. So, if you do computer verification, which is good, it's a very

strong tool, you have to have the computer do warnings that something's wrong. The independent verification – the key is it has to be somebody who's independent; it can be a second person. It can be running things on a second computer, if all the inputs are new, because the most likely errors, if you're doing checks of, say, a treatment plan, the most likely errors are in the input, not in what the computer does. And so, what you don't want to do is take the output from the first computer, for example in a brachytherapy case, you take the coordinates from the localization that you found in the first computer and enter them in the second computer. All you're doing there is checking that the computational algorithms are the same, and that's something you should have done during acceptance testing. The computers, as I said, can reduce errors tremendously, but what you've traded

off, is you've traded off types of errors, and you've opened new avenues, the avenues being that people depend on the computers and stop looking at what's going on.

“Redundancy is not useless duplication.” That's a quote from Robert Loevinger, who liked to say that, and it's really wise. It's ... redundancy looks at something twice, and in different ways, so that you can have different perspectives to try to find errors. On the other hand, useless duplication is not redundancy. If you just do something and just do it again, that is, put the input into the computer and do it again, that may not be appropriate redundancy. Yes, if you're checking your inputs and making sure that you've put things into the computer correctly, that is valuable, but if whatever you're doing isn't really starting over again and having new inputs, then it is just useless duplication, and it's not

an effective check. We have all sorts of levels of quality management that we do as physicists. We do acceptance testing, commissioning, periodic testing at various periods, and stuff on each patient. As we're organizing our quality management, we have to plan, of course, what falls into each of these categories. And that goes back to the question of periodicity that we talked about earlier. Let's see, given the time, I'll skip this. This is just talking about in radiotherapy, what are things important to check on. The plan – that's all in the handout. Let's look at misadministration's that have been reported to the Nuclear Regulatory Commission and the International Atomic Energy Agency. We did a study looking at misadministration's, and analyzing what happened. And some of the ... some of the commonalities can be useful. A lot of the misadministration's occurred because of

failure to consider the human performance in the design of equipment. That was actually the largest single problem. The equipment, per se, didn't fail, but the design led to people making mistakes. You could see how the design led people to just go down the primrose path without realizing that something was going wrong. The most common in brachytherapy was high-dose-rate brachytherapy a default distance that was built in to both the treatment planning computer and the treatment unit computer, that if it wasn't overridden, would, with just hitting the return key, give you a default distance, and it

wasn't the distance people wanted, but they ended up – because they most often used the default distance – they ended up doing that automatically. And a particular danger was those situations where the equipment suddenly malfunctioned, that it didn't work, and

suddenly the operators were forced to perform the functions that the computer normally did, and while they were ... they were functions that the people could and should be able to do, they weren't used to it any more. A final one is entry of data in terms, in units, usually, other than what people are commonly used to working with. This becomes a big problem, for example, as you switch from doing your treatment plans in millicuries to air-kerma strength. There's been a lot of errors involving that. Events tend to happen most with actions having the least available time. This is not a big surprise, and people have known this from industry for a long time, and we all know that, that when suddenly we're told we have to get something done right now, and you don't have enough time to think about what you're doing, that's when errors happen. The events are often followed by the

failure of people involved to detect that the situation is abnormal, even though, even though they had all sorts of inputs telling them something's wrong. A typical example is the SL-1 reactor, an army reactor in Idaho that blew up. There's a story, it was a good story, behind that one too. But right now we'll focus on the first person to know that that happened was a sentry at the gates to the camp three miles away because the radiation alarm went off in the booth, and a well trained and attentive sentry grabbed the telephone right away and called the appropriate person and said, "Are you the electrician? Come and fix my alarm. It's going off. Something's wrong with it." That's really typical. There's a gene in the human that leads us to say, "Oh, everything's really right. Any indication that it's wrong is wrong." Once identified, the responses often include actions

appropriate for normal conditions but inappropriate for the conditions of the event, as we're very well trained to do what we do, and we do it even when it's not right. Lack of training to the ... is ... I see I didn't make a real sentence out of this. Lack of training to the point that the persons involved understand the principles was a common thing, that is they were trained, they were trained in what to do, but they didn't get enough training to understand what was going on, so when things started falling apart, they didn't know what to do. They couldn't improvise. And lack of procedures led to the problems, that is, they got outside of the normal procedures. They got off the yellow brick road and just didn't know where to go from there. [Here the text jumps back to "right away.and called the appropriate person..." and pickes up below at "first nurse..."]New procedures were

a problem, or new persons joining a case in the middle, often presented a hazard. A common ... a common problem occurs in nursing units when they hand off a patient from one nurse to another, and the first nurse knew about something with the patient that didn't quite get into the chart or was a bit idiosyncratic and the new nurse doesn't really know it, and it wasn't passed along. That's a problem. New procedures, or an old procedure with new people joining in, or a procedure with old people who knew what they were doing in a new location – all of those. All of those were dangerous situations. Most events suffered from ineffectual verification procedures, that is, they had QA in place, but

it didn't work, either because people didn't do it, they didn't do it right, or it just wasn't designed so it corrected. Common ... common situations that have happened. First, errors

don't happen from single causes. We're very good at handling things when they just go wrong. We're not so good when things go wrong, and something else goes wrong at the same time. Then we're confused, and we don't know where to put our attention. Errors are surrounded, usually, by indicators that something's wrong, but they aren't noted. Distractions are very common in misadministration's. Distractive work environments are a danger. Rushing, as I said, this is ... this is a real problem, and it's often just due to the pressures that we've got to do a lot of other things, and lack of staffing – big common problem. Lack of communication is a common cause of the problems. Insufficient and appropriate staffing. Oh, gee, we've had that up there five or six times in this list. Erroneous data entry is common. Finally, where have we come through all this? When

establishing your quality management, you should use the risk assessment tools to determine what needs to be done and what needs to be protected. Remember, you can't do everything, so you have to prioritize. Organize the quality management steps by QA and QC so that you can sort this out on your error tree. Recognize that errors will take place. This is probably the most important one. Lucian Leape, who has been in the forefront of error reduction in medicine, says, "Our job is not to prevent errors, but to keep the errors from injuring patients." And that's what you've got to keep in mind as you design your quality management. Thank you.