Timothy Unruh: Hello. I'm Tim Unruh, program manager for the Department of Energy's Federal Energy Management Program. Thanks for joining us for First Thursday Seminars. If you've participated in previous seminars, welcome back. If this is your first time, you can access our earlier seminars at any time by visiting the FEMP website.

Energy expenditures in our country are more than $1.2 trillion. Working to reduce our dependence on foreign fuel, while also reducing our energy costs, should be a priority for us all. Today, we have the technology, practices, and progress to make energy efficiency cost effective. Now is the time for all of us to take action and make our operations as cost efficient as possible. With new developments at home and abroad, the president has challenged the federal government to lead by example and meet aggressive goals to reduce the use of carbon-based energy and increase renewable energy production.

FEMP is producing this training to help you reduce costs associated with energy. To this end, FEMP assists facilities managers by offering resources and training for cost-effective energy management practices and smart investment decisions. FEMP training can help you identify, plan, and implement sound financial and technical solutions for a broad range of energy projects, no matter how large or small.

We hope this seminar, and the others in this series, will help you reach your energy, water, and greenhouse gas reduction targets and reduce your energy costs. Working together, we can reduce the energy costs of facilities, improve the composition and fuel economy of fleets, employ the federal building stock as an example for innovative efficiency methods, and utilize new technologies to increase energy security. Moreover, we can accomplish these goals through the use of project funding tools, such as energy savings performance contracts and utility energy service contracts, to stimulate the economy, create jobs, and reduce costs.

I invite you to visit the FEMP website for additional resources, technical assistance, and guidance and to call on our customer service representatives and specialists to share their expertise. Enjoy today's session. Log on to our previous seminars, and be sure to register for the ones to follow.

Kathy Hyland: Hello. Welcome to the Federal Energy Management Program's First Thursday Seminars. I'm Kathy Hyland, and I will be your moderator today. This is the sixth course in this series, and we'll focus on energy efficiency in labs, data centers, and high-tech facilities. If you have not already printed a copy of the learner guide and handouts, you may do so now. You can do that by accessing the website on your screen, www.femp.energy.gov/firstthursday. You can also print them out after the seminar. The materials and a video archive of this presentation will be available 24 hours a day/7 days a week.

Now let's look at the objectives for our training today. It is our hope that after completing this training, you will be able to; one, discuss the potential benefits and cost-savings potential of improving energy efficiency in labs, data centers, and other high-tech facilities; two, suggest technologies for improving energy efficiency in data centers, such as IT equipment, software, air management, humidity control, and free and liquid cooling; three, suggest technologies for improving energy efficiency in labs and high-tech facilities, including fume hoods, ventilation and air changes, and HVAC systems; four, discuss ways to improve energy efficiency through improved operations and maintenance; five, discuss FEMP tools and resources to support energy efficiency in labs, data centers, and high-tech facilities.

Our format today is simple. There will be a presentation followed by a question-and-answer session, and we really encourage your questions. From time to time on your screen you will see an email address, a fax number, and a phone number so you can pose your questions. If you would like to speak to one of the instructors live, you can phone in your questions and someone can speak with you directly, so we really encourage you to do that.

We have two instructors today. Dale Sartor leads the applications team at Lawrence Berkeley National Laboratory, where he assists in the transfer of new and underutilized technology through project-focused multidisciplinary teams. He is a licensed mechanical engineer and a licensed building contractor with more than 35 years' experience in energy efficiency and renewable energy applications.

Dr. Paul Mathew is a staff scientist and deputy leader of the commercial building systems group at Lawrence Berkeley National Laboratory, where he conducts applied research and market transformation activities on energy use in buildings. His current work is focused on energy efficiency and green design for laboratories and data centers, as well as energy benchmarking tools for commercial buildings.

We also have with us live from Washington, D.C., at the Department of Energy, Will Lintner. And I will introduce him further later.

So without further ado, let me turn the presentation over to Dale.

Dale Sartor: Thank you, Kathy. So today we're going talk about some of the big energy hogs in buildings, specifically data centers and laboratories. And as this slide shows, the typical laboratory is three to eight times more energy intensive than office and other commercial buildings, and data centers can be ten to 100 times more energy intensive.

There are a number of drivers for saving energy in high-tech buildings. For example, the Energy Independence and Security Act of 2007 requires that there be an 18% reduction by 2011 and a 30% reduction by 2015, based on 2003 baselines. High-tech buildings or energy-hog buildings are not excluded; neither are their processes, from these requirements.

In addition, there are other federal requirements. EISA also requires that we do energy audits for water and energy, and benchmark our data centers and our laboratories. The administration has issued – or has initiated the Federal Data Center Consolidation Initiative, which aims to close down or decommission 800 federal data centers. An executive order, such as 13514, contain additional requirements for laboratories and data centers.

So, first, we're going to talk about energy efficiency in data centers. Data centers are energy intensive. Individual racks can utilize up to 25 kilowatts of power. There's a surging demand for data storage. A typical facility can be a megawatt, or all the way up to 20 or more megawatts. This all adds up to about 1.5% of the U.S. energy use. That's equivalent to more energy than the entire state of Tennessee, where we're today. The energy operating costs are usually the highest operating costs in a data center, and in fact the energy cost is approaching or exceeding the cost of the IT equipment.

The Lawrence Berkley National Laboratory or LBL feels the pain. You can see, at the beginning of the century, our energy use in our supercomputer center was small, under several hundred kilowatts. But our projections were that we'd see growth all the way up to 35 megawatts. To put this in perspective, our entire campus's energy use is only 13 megawatts, so this one building was going more than quadruple than our entire campus's energy use. Our director at that time, Dr. Chu, was not pleased with this prediction, and he negotiated with our computing scientists, and everybody agreed that our energy for the data center would be capped at 17 megawatts. But even at 17 megawatts, that will more than double our energy use at LBL.

Energy efficiency in a typical data center is actually less than 15%. We start out with 100 units of energy at the source, or at the power plant. About two thirds of the energy is lost in the generation. Then another couple of percent is lost in distribution. And of that energy that makes it to the actual data center, oftentimes half or more of the energy is used for cooling and conversion of power, or the power chain. So, again, less than 15 percent actually gets the computer load itself.

We can actually save energy, or there are energy efficiency opportunities in each one of those steps. Starting with generation, we can provide on-site generation. Moving to the power conversion and distribution, we can provide the power in higher voltages that saves energy. We can have energy-efficient uninterrupted power systems (UPS) and more efficient redundancy strategies. Moving over to the IT equipment or the server load, we can pick energy-efficient components. We can virtualize our servers so they are more highly utilized. We can use more efficient power supplies, and we can implement load management. Then, finally, on the cooling equipment, we can improve our air management, we can move to more efficient liquid cooling, and we can optimize the cooling plant.

The potential benefits of data center efficiencies are large. We can save 20 to 40% savings in a typical data center. But with new construction, on an aggressive strategy, we can save 50% or more. In addition, energy efficiency extends the life and the capacity of the infrastructure. But do we know if our system is good or bad? That is where benchmarking comes in. Benchmarking allows us to compare ourselves to our peers, and what we find is there is a large variation in performance. Benchmarking also helps us to identify the best practices, helps us select projects to address, and then allows us to track and measure performance over time.

Here is an example of two data centers. Data center one, approximately half the power going to the data center is actually used by the IT equipment, whereas in data center two, about two thirds of the power is used by the IT equipment. Now, if everything else were the same, we would consider data center two to be more energy efficient.

So that brings us to our first high-level metric, the power utilization effectiveness. Here, that equals the total data center power divided by the IT power. So one is the IT power, so you can see for all data centers, they all have the one. But now, whatis above the one, the blue portion of the graph, that is the power that is used for the infrastructure, for the cooling and the electrical distribution chain. And as you can see, there is a huge variation from the best data center in this data set being about 20% of the IT load going to infrastructure, whereas in the worst case it is more like 200 or 250%. So the PUE ranges, in this case, from about 1.2 to over three.

So how do we achieve those levels of performance? How can we get it down to the 1.2 level rather than the three level of PUE? That's what we are going to talk about, the technologies and best practices for data centers.

The first opportunity is in the IT equipment itself. We can improve the computational efficiency through, for example, specifying Energy Star servers. We can consolidate and virtualize our servers. That means bringing the utilization of the equipment up from, say, 5 or 10% to 60 to 80% utilization. We can better manage our energy storage. That means putting storage or putting data that we do not access very often in lower forms or less energy-intensive storage media. We can implement power management. That is turning off equipment or turning components of equipment off when they are not needed. And then each one of these opportunities has a multiplier effect. As I mentioned, if we assume the average PUE is two, that means for every kilowatt the IT equipment uses, there's another kilowatt used for the infrastructure. Anything we save on the IT side is another watt or kilowatt saved on the infrastructure side.

The next opportunity is to use information technology to save energy in information technology. Operators of data centers lack visibility into their data center environment. If we could provide the same level of monitoring and visualization to data center operators on their environment, as they now have in their IT environment, we would be much better off. We would be able to measure and track performance, and spot problems early before they become major.

LBL, for example, has installed a 700-point wireless monitoring system produced by SynapSense that monitors the temperature, humidity, pressure, and power of our data center. Now, with 700 sensors, it is quickly a problem of too much data and what do you do with the data. So, fortunately, the visualization tools that we have are getting much, much better. This is an example of four shots that in one picture shows us all 700 points. These are not a simulation. This is real data. The first graph on the lower right is the under-floor temperature. The graph on the – or the picture on the lower left is the temperatures at the floor level. At the upper right it is the temperatures midway. And at the top of the racks we see the temperatures in the upper left-hand corner.

The ideal situation would be to see consistent temperatures, all blue, under the floor; and then cool temperatures in the cold aisles and hot temperatures in the hot aisles. Not only can we get a snapshot instantly of 700 points, but we also have animation where we can see literally thousands, tens of thousands of data within a few seconds displayed on the screen.

We can also, at the LBL case study, there is a real-time PUE display. So on the left we can see the breakdown where our energy is going over time, and then on the right is the PUE in real time. So anytime there is a change, our operators instantly know whether it is good or bad, and they can track down what the cause was.

Our next opportunity in data centers has to do with the environmental condition. How many people have been in a cold data center? Well, I know I cannot hear your answer, but my assumption is most of you have been cold in a data center, and that is a sure indication that there is an opportunity for savings.

Fortunately, ASHRAE has worked with all the computer manufacturers to develop recommended conditions, and they have all agreed that data centers can be up to 81 degrees. In other words, a computer is happy in an environment that is uncomfortably warm for its occupants. So in reality, you should be warm in a data center, but just the opposite is often the case. Now, we just can't go into a data center and crank the temperature up. We have to deal with addressing the air management problems first.

So, generally, there is a lot more air circulated in a data center than is required, and air mixing and short circuiting leads to low temperature supplies and low Delta Ts. So the first step to air management is to use hot and cold aisles. Then, once we have oriented the computers and the racks such that there is hot and cold aisles, we want to improve the isolation of the hot and cold aisles. By doing so, we reduce the fan energy, we improve air-conditioning efficiency, and we increase the cooling capacity.

I am going to run through some examples of air management improvements, and we are using the monitoring system at LBL to show the impact. The first improvement has to do with blanking panels. It is very typical in a computer rack to have space between one server and another server. There may be a missing server. Air from the hot side can easily pass through the rack and mix with the cold air on the cold side, or in the cold aisle, thus raising the temperature of the cold aisle. In this case, we installed a 12-inch blanking panel in the server rack. And what you can see is, the temperature on the cool side at the top dropped 20 degrees.

Another big opportunity in air management has to do with too many permeable floor tiles. Oftentimes, we discover there is a hot spot at the top of a rack, and the solution to that is to replace the floor tiles with permeable tiles or to replace less permeable tiles with more permeable tiles. That might solve the problem in that rack, but it reduces the pressure under the floor, and now we have hot spots throughout the data center. So by optimizing the air flow and the floor tiles, we can increase the pressure, and we actually decrease the temperature at the tops of the racks. In addition, the capacity of the overall data center is increased.

Another tool that we use was enclosing the return air. We had an overhead plenum that was converted to a hot-air return. Grills – exhaust grills or registers were placed over the hot aisle to collect the hot air as it leaves the computers. It is transferred over to the CRAC units via the ceiling plenum. And then intakes were extended from the CRAC units, which stands for computer room air conditioners, to the overhead return. Thus, we have isolated the hot air.

The next step was to install air curtains. And here, again, we' are using these curtains to help isolate the cold from the hot. In the ideal situation, we now have – instead of supplying air at 45 to 55 degrees, we are supplying air at 70 to 80 degrees. But cold-aisle temperatures should be equivalent to the supply air temperature and should be uniform, because it is in essence a plenum in and of itself. Then the air goes through the servers. They pick up 20 to 30 degrees. It's collected in the hot aisle and returns to the CRAC units at 95 to 100 degrees rather than 60 to 70 degrees.

Okay, the next opportunity that we are going to explore is the use of free cooling. There are two options to that. One is outside air economizers, where we use outside air directly to cool our computers. Sometimes the configuration of the building does not allow for that. So in that case, we can use water-side economizers, where the water from the cooling tower is used directly to cool the computers, or indirectly to cool air that then cools the computers. But in either case, we are reducing the need for the chillers, and the ultimate goal is to get rid of chillers in data centers.

Sort of in conjunction with that water-side economizer or tower-side economizer is the emerging technology of liquid cooling. Now, it is kind of funny to say "emerging" because, historically, computers were cooled with liquid. But in the last 20 years or so, they have been mostly cooled by air. But we are seeing a trend back to cooling by liquid, because liquid is much more efficient than air as a transfer fluid, and it allows us to use that free cooling more often.

So at LBL we installed what is called rear-door coolers, or basically coils on the backs of our computer racks, that take the heat immediately off the backs of the servers and cool it down before the air even leaves – or returns to the CRAC unit. Now, the way this works is, most of the year – and, for us, all the year – we are cooling that water in the rear-door coolers only with cooling-tower water. The plumbing is set up in such a fashion that if the cooling tower water is not cold enough, we can actually reduce it a few more degrees using a chiller. But in both options, whether it's chiller assist or cooling tower only, the system is much more efficient than our original DX CRAC units.

The last HVAC opportunity I want to talk about is improving humidity control. Oftentimes, we inadvertently cool air in a data center because the air conditioners are actually extracting moisture from the air. That extracted moisture has to be replaced by the air-conditioning system, and it is often done very inefficiently. ASHRAE's established, just like temperature ranges, very wide humidity ranges that are much wider than what the CRAC units come factory set at. So allowing for that and utilizing that range is really good.

The third point is we want to eliminate the fighting of some CRAC units dehumidifying while others are humidifying. And here is a table of air conditioners or CRAC units at LBL, where you can see that three of the cooling units were – or, I am sorry, three of the units were cooling only; two were cooling and humidifying; and one was cooling and dehumidifying, so they were fighting one another. We simply disconnected the entire humidity-control system, and what you can see is, the result is our humidity levels went down 3%. Perfectly acceptable, but our CRAC power went down 28%, so a huge saving opportunity.

Now I am going to move to opportunities in the power chain. Typically at a data center, power comes in at a high-voltage AC. It is converted to DC and runs through the UPS system. Then it is converted back to AC and increased again in voltage, where it goes over to power distribution units that reduce the voltage down to 240 or 120. It goes into the server boxes, where that power is converted back to DC – from AC back to DC. It is increased in voltage, and then it is stepped down to a number of voltages. So you can see that the power goes through a lot of changes, from the entrance to the building to the components on the motherboard. And each one of those conversions is a waste of energy and a potential for energy efficiency.

Looking at just the power supplies in a server, you can see they are typically sized for 100%, you might buy a 400-watt power supply. But, typically, the actual use of the data center or the server is only 200, so you're operating at 50% load. But oftentimes, or almost all the time, we will put in two servers – I am sorry, two power supplies in each server. So instead of operating at the 50% mark, we are operating at the 25%. And as you can see, the efficiency of the power supply starts to drop off at 25%, very rapidly. So it is not uncommon to see the power supplies being a huge energy waste at the point that we are actually operating them, which is in the 25% or less range.

We have a program now called 80-plus, which means the energy efficiency of the power supply is 80% all the way down to 25% load, so that is a really good program. Next?

The UPS system, similarly, has a curve, whereas if we are using redundant operation or if we are not using redundancy, the UPS system is fairly efficient, well above 90%. But then in a redundant situation at low load, you can see how the efficiency considerably drops off. And in fact, one data center in this sample said – this is actually major data. One of the data centers, over half the energy was lost in the UPS system before it even went into the data center. And I am sorry to report that that was a federal data center, agency to remain unnamed. Next.

So the last opportunity I want to talk about is operations and maintenance. There is a lot to be saved here. The most important step is to get IT and facilities staff talking to one another. We should use lifecycle-cost or total-cost-of-ownership models when we are making our purchase decisions and design decisions. We need to document the design intent, and train our staff accordingly. For example, sometimes we will go into a data center where there are permeable tiles that have been put in the hot aisle because the operators thought it was too hot. Well, the hot aisles are supposed to be hot. So everybody needs to understand the design concepts.

Then we want to benchmark and track performance. That leads us to better savings or better operation. And lastly, we want to commission and keep our equipment finely tuned. Thank you.

Kathy Hyland: We'll be back to Dale shortly, but now Paul Mathew will talk about some of the tools for data centers, including the use of the DC profiling tool.

Paul Mathew: Thank you, Kathy. So there are many different resources for implementing efficiency in data centers, especially if you are a data center owner, operator, if you are an architect or an engineer. And what I'm going to do is to briefly cover some of the DOE tools in this particular arena.

So listed here are a selection of these tools that you can find on the FEMP website. The first is a quick-start guide that can be used for you. If you have never really touched your – looked at the efficiency of a data center, that is probably a good place to start. We have a suite of analysis tools, and I will be actually giving you a demonstration of one of those shortly. There is a range of best-practice guides, many of which cover the practices that Dale was just covering in his part of the presentation.

We also have a data center programming guide, so that would be for those that are starting – actually looking at constructing a new data center, where you would want to understand what the requirements might be from an efficiency standpoint, and ensure that those are included in the request for proposals (RFPs) and so on. And, finally, we have a set of technology case study bulletins. What is interesting about these resources is that not only do they describe the technology itself and how it was implemented in the data center, but they also describe some of the nontechnical challenges with implementing some of these efficiency measures, and how these different agencies were able to overcome those challenges in deploying these efficiency measures.

So with that, I am going to talk a little bit about one of these particular tools, which is the data center profiling tool, also called DC Pro. It's part of the DC Pro tools suite. The intent here is to look at the overall efficiency of your data center, as well as some end-use breakouts, and potential areas for efficiency improvements. And that will cover air management, electrical systems, IT equipment, and cooling as well.

I should emphasize that the intent here is not that this tool actually does a full-scale investment-grade audit for you; o f course not. But it is meant to get you started. It gives you a list of potential efficiency opportunities, which then allows you to go to a consulting firm or your own facilities department to then look at what the next steps might be in terms of actually implementing those measures. So it is really what it says it is; just a profiling tool.

So with that, I am going to break and go to a demo, and hopefully Murphy's Law will not apply and the demo will work as we planned. So this is the homepage of the DC Pro profiling tool – whoops, I mistyped that. And, again, it's free to use, and anyone can use it. Anyone can sign up for an account online.

And so once you log in, it comes to this first screen, which – where you provide your basic – what they call case information. "Case" is just a term for – each data center, for instance, would be a separate case. Here I have – this was a demo case user, and so I already have several cases loaded in, so I am just going to select one of the ones that we already have there to give you a demo of. And here you can see, once you have loaded that in, it shows you some of the initial characteristics of the data center; the floor area, the type of the data center, the tier level of the data center, and so on. And so that is the first step. Once you have entered some of that information, you hit the "save and continue" button, and now you come to step two.

This is where you begin to describe the actual characteristics of your data center. Again, it is at a fairly high level. And you will notice across the top we have tabs for the different areas of your data center. Your energy management practices, starting on the way left here. Then we look at IT equipment. And in the IT equipment, for instance, it asks you some simple questions such as "Do you measure and track your server utilization?", "Do you have a process for identifying abandoned or unused servers?", "Are you using virtualization?" and so on. And most of the answers, as you will see here, are in simple dropdown menus, or they are radial buttons, or in some cases there may be some text-entry fields as well. Again, very few of these are actually required answers. If you do not know the answer, you can just leave it blank and move on, and come back later once you have been able to obtain the answer to some of these questions.

I am just going to breeze through a couple of other tabs here just to give you a flavor for the kinds of questions we ask. Under "environmental conditions," for example, we ask you what your typical average supply temperature is. We ask you whether you have humidification controls. Dale spoke briefly about the importance of humidification controls in terms of being an efficiency opportunity. And then whether you have procedures and personnel for grounding cable equipment to prevent ESD and so on.

Perhaps I'll look at one more tab here before we move on. Under "air management," we ask questions related to essentially how you manage air within your data center, and the air supply paths. We look at questions related to, for instance, cable penetrations and whether those are sealed; about your IT equipment and whether it is organized in rows, and how well it is rganized in rows, and how tightly you are managing the rows, and so on. So, again, most of these questions are fairly straightforward. You should be able to answer them right on the spot there, but, again, you can always come back later and fill in the answers if you do not have them right on hand.

Well, maybe we will just do one more tab here. Under "IT equipment power chain," we ask questions related to UPS technology types, also about your PDUs and the types of PDUs, and so on.

So once you have answered a selection of these questions, you then go to the default breakout screen. And what the tool does for you here is that it takes the answers you have provided, and then it creates a default breakout of the energy use in your data center. Now, obviously, at this point you have not provided any actual energy-use data for your data center, but it is creating a default based on some of those answers. And in this case, for instance, it shows that about 61% of your total data center energy use is probably going to IT equipment.

So once you have done with that, you can move to step three. And in step three, it asks you – this is an optional step. It asks you to enter data related to your actual procedures that are going on in your data center. For example, if your data center is primarily for storage, you might enter the total amount of storage that you have. This is so that you can generate productivity metrics for your data center, such as BTUs per terabyte of storage, and so on. Again, this may not be applicable to all data centers, so if it is not applicable to yours, you can skip it and move on to the next step.

And that would be step four. And in step four you enter your energy use information. This, again, is fairly straightforward. If you have used – or if you have EIS systems, you will see that this information is, again, typically the kind of information you would enter in an EIS system. So we have four different fuel types here. There is electricity, fuel – that would be natural gas, fuel oil, and so on – steam, and chilled water. And chilled water, in this case, refers to supplied district chilled water, not chilled water that you are maybe actually generating at your data center itself. So the data itself that you enter is essentially annual use. If you have a quarterly use, a monthly use, and you enter that, we will essentially take that and annualize it. So it is best to actually enter your annual use.

So once you have entered your overall fuel consumption and electricity consumption, you then have the opportunity to provide a breakout for each of those meters that you entered. Here, again, the idea is that you provide as much information as you have. If you have a sense, for instance, how much energy is going to your IT equipment, you can indicate that here on the screen. If not, you can leave it blank and move on. And, again, the tool will use whatever inputs you have provided, in order to be able to give you some answers on the results page. Again, you can always come back, and if you have refined your numbers, you can refine them – you can enter them here and so on.

So once you have entered those data, you then come to this screen, which summarizes for you the breakout of your energy use in your data center. And you will notice, after we have entered some data, this now reflects an adjusted distribution of your data center energy use. And in this case, it shows about 57% of your total energy is going to IT, and the remainder is actually used for infrastructure.

Now, I should mention and maybe interject here that this is the current version of the tool. That is, again, available on the website. In about a month, we are going to be releasing a new version with some enhancements on usability and also to ensure that the new version uses PUE rather than DCiE. DCiE is actually – stands for data center infrastructure efficiency. It is actually just the inverse of PUE. And when we first launched this tool, it was considered a popular metric, but now the industry has really moved much more towards PUE, so we are going to change that and represent it as PUE. But, again, one number is just the inverse of the other.

All right, so once you have entered all those data, you then come to the results page. On the results page, as you can see, first it just presents back to you the annual energy use data that you provided. Then, under potential energy savings, it presents you your current energy use, both in terms of site and source energy, and then it also presents to you your potential energy use – that is what essentially the best practice would be for your data center. And the difference between the two essentially represents your potential savings. Those same data are then represented also graphically, below.

And, again, we should interject here that the potential savings are really representing almost an idealistic case for your particular climate. It does not mean that you could actually get there for your particular data center. There are obviously site-specific constraints that may preclude you from doing that, but it is an idea to give you a sense of what an aspirational goal might be for your particular data center.

And further down, at the bottom here, we have a benchmarking chart that shows you where your data center currently stands relative to a peer group of data centers that we had measured earlier and then where you potentially could be. So your current DCiE is about 0.57, and you could be at about 0.9 for the climate zone that you are in.

And the last part of the results section is that we provide you a list of potential energy savings measures. What we did here is that the tool has built into it a set of about 80 measures, 80 different energy efficiency measures. And based on the answers you provided, we then cull through that list and pull out the ones that are likely to be most effective for your particular data center and present those back to you. You can then make a PDF printout of this, and that essentially acts as a starting point for the next step for you, which is, again, to call in a consultant or to do a more detailed analysis. And that, again, would help you proceed towards implementing these measures within your data center. And I should say, some of these measures may not even require further detailed analysis. They may be very simple, extremely low or no cost, and you could just go ahead and do them. But some, of course, will require much more detailed analysis.

So with that, I will go back to our presentation here and talk about a couple of other resources. One is that we have a couple of training opportunities. ASHRAE and DOE have an awareness training; this is a one-day training that is targeted at data center and facility operators. We also have a certified practitioner program. There is also a three-day generalist workshop, and there are two options. The first option, there's a training certificate track, which has no pre-qualifications. You attend the training and you get a training certificate. And the second option is where you do the training, and then you actually take an exam. And if you pass the exam, you then get to be a DCEP certificate, essentially considered a data center certified energy practitioner.

The target audience for this training is data center personnel, consultants, and service providers. We also have specialist tracks for electrical and air management, and soon – or, rather, for air management and HVAC, and soon we will have it for all of these areas that we have listed here.

And, finally, before I hand it back to Kathy, I would like to say that – I would just like to point out some of these websites that have a lot of these resources and additional information. Of course, the FEMP website itself. Lawrence Berkeley Lab's website on data centers has, again, a lot of resources, including some that were developed specifically for the California Energy Commission, but are resources that could be applied much more broadly nationally as well. Energy Star's website has good information on server efficiencies, as well as the Portfolio Manager tool that can be used to benchmark your data center. And DOE's Industrial Technologies Program has a Save Energy Now website that also has a rich set of resources to help you get started with your data centers.

So with that, Kathy, I'll turn it back to you.

Kathy Hyland: Okay. Now let us hear from Will Lintner, and then we will move to the last section. Will Lintner is a team member of the U.S. Department of Energy, Federal Energy Management Program, and he is responsible for facilitating the use of new energy efficiency, water, and renewable energy technologies and sustainable best practices within federal facilities, including laboratories and data centers. ToWill.

Will Lintner: Federal research laboratories and data centers are a critical infrastructure needed to continue advances in science and technology, essential to American innovation and security. When compared to other federal facilities, these facilities use much larger amounts of energy. As you are learning in this seminar, there is a great opportunity to reduce energy use, because new technical resources are making energy-smart solutions readily available.

Hello, I am Will Lintner, data center/Labs21 project manager for the Department of Energy's Federal Energy Management Program. Operating efficient labs, data centers, and other energy-intensive facilities is good for the economy and the environment, and required to meet federal mandates, including Section 432 of EISA and Executive Order 13514. Many of these facilities have been overlooked in the traditional discovery process for energy-savings opportunities. This is usually because of a lack of awareness and a perception that the technical or organizational challenges are too difficult.

The key point of this seminar is that there are many opportunities for improvement that are cost effective. Some of the opportunities are fairly simple to implement, with little cost. Your agency, too, can realize a quick payback on efficiency upgrades by taking advantage of FEMP training, tools, and resources. FEMP's comprehensive *Best Practices Guide to Energy-Efficient Data Center Design*, our *Benchmarking Guide*, and other publications will help you improve, measure, and benchmark data center performance. For instance, FEMP has been creating technology case studies to provide you with information on the performance of energy-efficient technology. We collaborate with others within the Department of Energy, throughout the government, and in the private sector to develop and provide access to awareness training, professional certification programs, and assessment tools.

To accelerate best practices and technology deployment, FEMP participates in interagency groups, such as the Federal Partnership for Green Data Centers, a forum for agencies to collaborate and share lessons learned. FEMP also participates in the Federal Data Center Consolidation Task Force to promote a closer cooperation between the facilities and IT departments and to provide top-level IT professionals with technical assistance as agencies consolidate and transform their IT operations. We invite you to learn more about these and other groups and to take advantage of numerous resources by visiting our website.

FEMP also provides a valuable collection of resources for those charged with increasing efficient operations at federal laboratories. Visit the FEMP Labs21 portal to see how FEMP actively participates with EPA and others in Laboratories for the 21st Century, a partnership that promotes high-performance sustainable building practices.

The Labs21 tool kit is an indispensible reference that includes the Design Process Manual, a tool that provides step-by-step guidance of the design process for a high-performance laboratory, leveraging many other Labs21 tools, such as case studies and technical bulletins. We have also developed a Web-based benchmarking tool to help agencies compare their laboratory's energy use with similar laboratories, in total and at the system level.

The Labs21 annual conference, sponsored by the International Institute for Sustainable Laboratories, provides an excellent opportunity to network with colleagues, take advantage, and see demos of the latest systems, technologies, and equipment. Each year, the conference features workshops, technical presentations from industry-leading architects and engineers, exhibits, and local tours of exemplary sustainable labs.

Reducing the energy and environmental footprint of federal labs, data centers, and high-tech facilities is readily possible and can lead to large savings. Others in the public and private sector are realizing energy and cost savings from these opportunities, and you can too.

Thank you for joining us today. I look forward to answering your questions later in this seminar.

Dale Sartor: Thank you, Will. Now we are going to turn to the specific opportunities available in laboratories.

As I said, laboratories are energy intensive, three to eight times the energy intensive as an office. And the potential for savings is also large, 30 to 50% over standard practice. In fact, at LBL in the 1990s we were able to reduce our energy use per square foot over our entire campus by 40%. That yields a triple bottom line, including a reduction in lifecycle cost, improved workplace quality and safety, and reduced environmental impact.

This slide, which depicts our benchmarking tool that both Will mentioned and Paul will explain in more detail, shows the high variability of energy use in laboratories, but also shows that they are all significantly higher than the typical office environment. Unfortunately, just like with data centers, a simple EUI isn't sufficient to tell whether your laboratory is good or bad. So throughout my presentation, I will try to talk about some of the metrics that apply to some of the best practices.

The laboratory energy use is dominated by HVAC (heating, ventilating, and air conditioning). The ventilation system is the largest component of the energy consumption. And in fact, in many labs, 10 to 20% improvement in the ventilation system is equivalent to the entire energy used by the lighting system. That is not to say that we should not take advantage of lighting efficiency opportunities, but we should focus on where the big hits are.

So we are going to look at five of the biggest opportunities in laboratories today. The first is scrutinizing the air changes, which optimizes the ventilation rates. Second is taming the hoods. Third is dropping the pressure drop, to lower the pressure drop in the HVAC designs. The fourth is to get real with plug load so that we right-size our HVAC equipment. And fifth is just say no to reheat, which minimizes the simultaneous heating and cooling.

The first one, scrutinizing the air changes. Air changes have a large peak and total-cost impact on laboratories. And we should not assume that air changes are driven by thermal loads. Oftentimes, that is what our engineers will tell us, and, frankly, it is just not true. So what do you use as a minimum air change rate? If your answer is "over six," you probably have a big opportunity, and you should quiz those that set these standards as to why the air change rate is what it is. And if the answer does not make sense, ask again. Ask why, why, why. If the answer is ten or more air changes – or, I am sorry, if the answer, again, is more than six, find out, well, when is ten or more air changes safe and six air changes or less is not?

Some of the options that can be considered in terms of air changes is, first, instead of looking at air changes, look at CFM per square foot, because oftentimes the air change requirements in a laboratory are focused on dealing with a chemical spill, and that is going be more related to square footage than it is volume. The next option is to consider a panic switch. In this method, you have one air change rate under normal circumstances, but as the occupant has a hazard that needs containment or needs to be dispelled, as they leave the lab, which is the safe thing to do, they would hit a panic switch that would increase the ventilation rate to clear the hazard.

The third option is to cascade air – in other words, use air more than once – to go from cleaner spaces to dirtier spaces. The next option is to set back the air change rate when the lab is not used, or to use demand-controlled ventilation, which monitors the hazards and odors in a laboratory and adjusts the ventilation rate only when there is a hazard present. The last option is called control banding, which in essence acknowledges the fact that different labs may have different needs. And instead of just having one-air-change-rate-fits-all, you would specifically pick the air change that relates to the hazard.

The bottom line is, the ventilation effectiveness is more dependent on laboratory and HVAC design than it is the air change rate. In fact, a high air change rate can actually have a negative impact on containment devices such as our fume hoods. Speaking of fume hoods that's our second big opportunity; to tame the hoods. The typical fume hood drives – or is responsible for consuming the equivalent of three houses.

What can we do with hoods? The first step is to scrutinize the number and size of hoods. And in new construction, ideally, you would design a system that allows it to be easily added and – hoods to be easily added and subtracted, so you would not have to start off with an excessive number of hoods. You would add some as you need them, and you would subtract them as you do not need them. It is very common to go into labs where they are starved for air because so many hoods have been added, yet there is additional laboratories that want or need additional hoods. But then as you walk around, some of the existing or many of the existing hoods in the labs are being underutilized or utilized improperly – say, for example, to store a lunch.

The second option is to restrict the sash openings, because the energy draw of a hood is typically dependent on what they call the sash opening, or the face velocity is set for 100 feet per minute. So the smaller the opening, the lower the airflow that we need, so restricting the sash opening is a tool.

We can also have unoccupied set points, occupied and unoccupied modes, where the airflow rate through the hood goes down when the lab is unoccupied. We can use variable air volume to adjust the exhaust so that the face velocity maintains at a constant 100 feet per minute. So as we close the sash, the energy use of the fume hood goes down, or the airflow through the fume hood goes down. If we can not depend on the occupants to close the sash, we can use automatic sash closures. And then we can also consider high-performance hoods. There is a whole set of new hoods coming out on the market that both reduce energy use as well as improve safety.

LBL has a fume hood calculator at this website that allows you to compare options and assess the performance of hoods or the cost of hoods.

The next opportunity is reducing the pressure drop. Up to one half of the HVAC energy use goes to fans. So the question is, how low can we go? Simple changes like increasing the duct size from 20 inches to 24 inches yields big savings, going from three-quarter horsepower to one-quarter horsepower.

Likewise, at the coils and filters, we can go from conventional design of 500 feet per minute, with a pressure drop of 0.8 inches, to just splitting the coil and doubling its size – not its volume, but its size, face area, reducing the face velocity to 250 feet per minute and reducing the pressure drop to 0.2 inches. So in this case, we have reduced the energy 75%. We have made our fans smaller and lower cost. We have increased the filter life, and the system runs quieter. An important point here is not to use rules of thumb that were developed for office design, because the ventilation systems in laboratories are much more highly utilized 24/7.

Here is an example of the cost impact associated with pressure drop or reducing pressure drop. This is a set of class five cleanrooms, so they are all providing the same level of service, but you can see there is a $400,000.00 difference between the worst and the best. That's $20.00 a square feet; these were normalized to a 20,000-square-foot facility. Next.

The Labs21 program has conveniently broken down each device in the system and created "standard," "good," and "better" metrics. The bottom line here is, in watts per CFM, you can see there is a factor of three difference from the best to the worst. Overlaying "good," "better," and "best," or "standard," "good," and "better," with our benchmark data, you can see that we are on target with actual performance. In fact many, many laboratories use more than even the standard, so there is a lot of opportunities for saving in our ventilation systems.

The fourth opportunity is getting real with plug loads. By correctly estimating the use of our processes in our laboratories, we can save capital and operating cost. And the key to doing this is by measuring actual loads in similar labs and then designing for good part-load efficiency, because you always want a safety factor; you always want a plan for flexibility down the line. So making sure that our designs operate very efficiently at low part loads is critical. The plug load diversity in a typical lab, which is high, increases the amount of reheat, and I am going to talk about reheat in more detail in a minute.

This slide depicts actual test data from a laboratory building at UC Davis, and it compares the design values to the actual measured values once the laboratory was built. And as you can see, there is a large discrepancy between the two, a lot of overdesign in the design phase. Next.

UC Davis is not unique. Sandia's PETL Lab was designed for 6 watts a square foot, and actual data showed that it used 1.8 watts average and 2.7 watts peak. The Fred Hutch Cancer Research Center in phase one was designed for 15 to 30 watts per square foot, but in phase two, after finding that the actual load in phase one was so low, they reduced the design to 8 watts a square foot. A Pharmacia lab was designed for 12 watts per square foot, and it was metered; data shows that it only used 2.7. These aren't isolated cases. These are extremely typical, and the only way we can deal with it is to design based on measured data of similar labs.

At LBNL's Molecular Foundry lab, we were able to reduce the first cost of a new lab by $2.5 million, by redesigning the HVAC system based on actual plug loads from measured data of similar laboratories. This freed up 4% of the construction cost, which allowed us to employ or deploy additional sustainability measures, so we were able to achieve a LEED Gold for that laboratory.

The last opportunity I want to talk about is just saying no to reheat, or simultaneous heating and cooling. There is only one lab in a laboratory building that the energy use or plug load far exceeds the rest, and oftentimes exceeds what the designers were anticipating. That lab drives the air-conditioning system to supply a very cold supply temperature, and then all the other labs have to have the air reheated. It is not uncommon for the boiler to be the largest load on the cooling system or on the chiller. VAV helps this situation but does not eliminate the problem.

There is a number of alternatives to eliminate or minimize reheat. The key is separating the ventilation systems from the cooling system. So if we stop having to cool the outside air to cool the inside, we actually can significantly reduce the size of the ventilation system and the chiller plant, so that when all the costs are considered, the cost of providing alternatives to reheat actually can be lower on a first-cost basis.

While these five big hits all provide significant opportunities, there are many other opportunities in laboratories, such as using energy recovery devices. And Paul Mathew is going to describe some of the resources that will help you identify all the opportunities.

Kathy Hyland: Now, just to introduce Paul Mathew again, he is going to talk about some of the Labs21 tools that are available to you.

Paul A. Mathew: Thank you, Kathy and Dale. Again, I just wanted to share some of the resources that we have in the Labs21 program. There are many resources out there. Time is going to prevent us from giving you a full overview of all of them, but we will just highlight a few.

Will, in his video, already mentioned that the Labs21 program is a great resource. It has been around for about a dozen years almost now and can provide you with technical assistance. There is also a toolkit, and we have training opportunities as well. And it is really an organization dedicated to providing – to making laboratories more efficient, not just in the federal sector, but more broadly across the country.

This slide summarizes the toolkit itself. Again, Will briefly touched upon it. Maybe a couple of other tools I would want to point out are the Environmental Performance Criteria, which is a bit like a LEED for laboratories, essentially, that we developed in response to many people in the laboratory community asking us to develop something that works better than the LEED for most other commercial buildings.

So with that, I would like to briefly cover two tools in that toolkit. One is what we call LEEP, the Laboratory Energy Efficiency Profiler. It is very similar to the tool for data centers where we profiled the potential efficiency opportunities in data centers. Here, again, we want to profile the potential efficiency opportunities in laboratories.

So, again, I am going to go to a demo here of the tool. This tool is, again, available – you can access it through the FEMP website or the Labs21 website. It is essentially leep.lbl.gov. And, again, it is free to use by anyone, and you can just request a user account here and we will send you a username and password that you can then reset.

So once you log in, the first question we ask you is whether you benchmark your facility, because often, again, it is good to start with just benchmarking, especially at the whole-building level, so that you understand that your facility may in fact have a big opportunity. LEEP, then, is the next step to identify specific opportunities within your facility. So we are going to continue with LEEP.

And the first screen – there are essentially four screens here. It's a fairly simple tool. The first screen shows you the list of facilities that you have in here. We, incidentally, applied this tool on a pilot basis to four USDA laboratory buildings, all located on the west coast. And they were very happy with the results because I think, fairly quickly, with about two hours of effort, two to three hours of effort, they were able to get a list of potential efficiency measures. And most of the data, again, that are required by the tool, most facility managers will have that pretty much off the top of their head, or they can go back and gather more information.

So I am just going to select one of these facilities. This is the USDA facility, actually, in Davis. And the next screen simply asks you the scope of your assessment. In many cases, you may not be interested, for example, at looking at lighting or at process loads, in which case you can just uncheck those boxes. We are going to look at all of them. And the screen right after that is essentially the core of the tool. That is where you begin to enter all your data about the systems themselves.

Under the "general" tab here, we ask about the facility and where it is located and its size and so on. Then, for each of the system areas on ventilation – again, we ask questions related to the total supply airflow, CFM, the efficiency of the fan loaders, and so on. Here again, most of the inputs are simple dropdowns or they are simple text-entry boxes, or they may be radial buttons as well. So, again, I am not going to go through each of those. I just want to give you a flavor for the kinds of inputs we are asking.

Under "heating and cooling," again, we ask whether you have a district system or if it is a dedicated system on site. We ask about control systems. Then under "process and plug loads," we ask about refrigerators and freezers and how they are – what their efficiency is, how they are configured, such as whether they are in dedicated rooms, and so on. And then, finally, we cover lighting as well. As Dale mentioned, lighting is not necessarily a significant portion of lab energy use. But nonetheless, there may be good efficiency opportunities there as well.

So once you have entered all these data, you hit the "save and continue" button, and it will bring you to this page. This is the last screen of the tool, and it has a list of potential efficiency actions. So those are listed here on the first column, "potential efficiency actions," and then for each action, we provide four pieces of information. The first is the relevance of that action. It could be high, or likely, or it is already efficient and therefore it is not relevant in your particular facility.

The next is the impact of that, whether it is a big, a medium, or a small impact. And you realize that it is possible for an action to have high relevance but low impact. Again, lighting would be an example where it may be highly relevant because you have inefficient lights, but because it is a small part of total load, its overall impact on your facility is going to be a little low. We then provide just some qualitative indication of the cost effectiveness of those particular measures, and then more information about what the next steps might be and where you can go for additional information to implement that particular measure.

So, again, as in the case of the data center profiling tool, this gives you a first cut. Some of these measures, again, are very low cost or no cost; you can probably just go ahead and do them. Others will require much more detailed analysis. Obviously, a chiller retrofit, for example, is not something you just then head out and do; you wanna do a much more detailed assessment before you proceed on that. So that is the LEEP tool.

The next tool I wanted to cover, that is the last tool, is the benchmarking tool. Again, Dale covered parts of this in his presentation, or he alluded to it. Here, again, we are first going to log in. And, again, this is freely available to anyone. You can just get an ID from the site itself. Then next, you are going to select which facility you wanna benchmark. And, of course, if you have not entered one before, then you can just enter a new one and it will create one for you.

So once you get in, the main screen, again, is where you enter the data about your laboratory building, about the type, the use, the number of buildings, the lab area, the gross area, and so on, the number of fume hoods. And then it asks you for your whole-building energy use data. This is, again, all the utilities that you have that are supplied to your laboratory building; not just electricity and natural gas, but also, for example, district chilled water or district hot water, if you have that.

If you only have estimated values, that is fine. Just indicate it accordingly, and you will see that those are characterized separately on the results page. You can then enter system-level data if you have those and if you want to benchmark at the system level. Again, none of these are required, but they are there for you in case you want to analyze system-level data. So, again, that covers ventilation, cooling, lighting, and process and plug.

So once you have entered all those data, there is the final screen that allows you to review your data, make sure that it is correct. And if it is not, of course, you can go back and edit it. And then we are ready to do some benchmarking analysis. So here, again, there are – we allow you to benchmark with a range of different metrics, so the first step for you to do is to select which metric you want to use to benchmark. It could be at the whole-building level, and it could be at the system level, with ventilation and lighting and cooling and so on.

So we are going to select – for starters, let us select at the whole-building level, and we are going to look at total source energy. And then you can filter the data set that you want to use to compare yourself against. You can filter that based on lab-area-to-gross-area ratio, occupancy hours, lab type, lab use, climate zone, and whether you want to exclude estimated data. So for starters it is kind of good to just leave in a fairly large data set and then filter the data set down just to make sure you have enough data to begin with.

So I am going to maybe select – or the lab we had input, the demo lab was about 0.4, so I am just going to select a range of 0.2 to 0.6 for lab-area-to-gross-area ratio. I am going to assume standard occupancy hours. It was a mixed chemical and biological lab, so I am going to select all those lab types. And it was in Virginia, so I am going to select Climate Zone 4A, and there is a little handy map of climate zones here that you can see. And then once you have done that, you hit the "continue" button and you get your results.

So these results show you the actual benchmark data. So this is, again, total gross BTUs per square foot per year. That is shown on the left-side vertical axis, and the dark green bars represent the peer data. Your own buildings are in light green, so in this case we have our one building here. And the red line across here shows you the average value for the peer group. So you can see we are significantly above the average, which indicates that there is probably a good efficiency opportunity here – above average meaning, of course, higher is worse here because you consume more energy per square foot.

Below the chart, we have the statistical information itself. It summarizes for you the filters that you use; the minimum, the average, the maximum of your peer group. And in tabular form, you can see the actual data for the peer group as well. All the peer group data are anonymized. You really cannot identify the buildings themselves. So if that is an important consideration to you, we just want to make sure that you understand that.

So I will just show you one more metric. Again, one of my favorites is ventilation watts per CFM, and for this one, let us expand the data set a little bit. And you will see here – this was representing the chart, again, that Dale had shown with peak watts per CFM. And our building, again, in light purple here, is just about average of the peer group in terms of watts per CFM, so it probably indicates that we can get a little more efficient.

And, again, one thing I perhaps neglected to mention was that if data are estimated, they have a hash mark in here, so that you can quickly identify estimated versus actual measured data in the data set.

So with that, I think, again, I will close. There are many other tools. We would encourage you to go to the website and look at these additional tools. They are a very rich data set, as I mentioned, that have been developed over several years. And so, finally, we should leave you with some contact information. And that is here for both Will, who is the program manager at FEMP for labs and data centers, and also Dale and my contact information. Thank you.

Kathy Hyland: Okay, thank you, Dale, Paul, and Will. And at this time, we will be happy to answer your questions, and so please give us a call so you can speak directly with one of our instructors or with Will in Washington, D.C. I have some questions that have already been posed, so let me start with those, and the first one is directed towards Dale. Dale, how will cloud computing affect data center energy use in the future? And a second question to that is: What impacts will facility managers see as a result of this change?

Dale Sartor: That is a good question. Cloud computing and consolidation actually may not save any energy, because it just depends on – it is just moving the use of energy from one location to another. But on the other hand, consolidation and going to the cloud is going to tend to facilitate higher levels of efficiency, because the data centers will be larger. They will be more professionally managed. Sometimes we will be going to the private cloud – or, I am sorry, the public cloud – going to the private sector to buy cloud services, where the cost of providing the service is very key, and variable costs like energy become the difference between making money and losing money. So in those situations, energy efficiency is likely to be more important.

For the energy manager, obviously, if you are the site that is losing the data center, your job got easier; if you are the site that is being consolidated to, the importance – just like at LBL, where we are seeing the data center contributing to two or three times more energy than our entire site – those consolidated centers are going to be very energy intensive. And it is going to put a lot more pressure on the energy manager to do something regarding energy efficiency at the consolidated centers.

Kathy Hyland: We have a second question.

Audience Member: Yes. Thank you very much. I enjoyed your seminar, and I was glad to see a lot of those topics, a lot of the – most all of these techniques we use ourselves in our designs. I wanted to ask: one, in a federal facility, whether it be a data center or laboratory, what type of return on investment or payback period does a facility entertain when they look at energy efficiency retrofit?

Kathy Hyland: Will, would you like to take that one?

Will Lintner: Okay, I can address that. What we generally have as a limit or a guideline is a ten-year payback on an energy conservation measure. That was actually put into one of the early laws. I think it was the EPAct 1992, actually put that into the statute.

Audience Member: Very good. I had a second question. I believe it was maybe about a year ago, I was looking for some different types of hoods, more energy-efficient hoods, and I came across a hood called the Berkeley hood, developed there at Berkeley. Do you know anything about it?

Dale Sartor: Yeah, I can take that question. Yes, I mentioned the last bullet on the hood options was to consider high-performance hoods. And the Berkeley hood was developed a number of years ago. It uses a push-pull concept. It's easier to blow out a candle than to suck out a candle. And, frankly, the hood has not been commercialized in the U.S. It was funded by State of California Resources – and California is the only state in the country that mandates a face velocity rather than containment of the hood.

But what the Berkeley hood did do was it stimulated a whole rift of additional designs for high-performance hoods. So there are many other ideas and concepts on the market, commercialized, that in essence do the same thing, which is improve safety at the same time of reducing energy use.

Audience Member: Well, thank you very much. I've enjoyed this.

Kathy Hyland: We have another one, and I am going to direct this one towards Paul. It says: Researchers and scientists often seem totally focused on their academic work. How do we get them on board and participating in these energy savings programs?

Paul A. Mathew: Well, yes, I mean, that is obviously a concern in laboratories – well, it is either a concern or an opportunity, I guess. I mean, in laboratories, for example, users are going to be critical to the success of measures, such as fume and sash management. And so engaging them up front is quite critical. In terms of how to engage them, one way – we've seen it done different ways.

One way, actually, is to actually find a champion. Try to find that one faculty member or that one scientist that may be really interested in energy efficiency or in environmental sustainability. Get them on board; get them to be your ally. They then act as an advocate to their peers. So that, I think, is a key element. In fact, we know one case where this faculty member – this was, obviously, a university, not a federal lab. But they got really engaged, and they themselves went around and, for instance, did a hood count, looked at what could be turned off, and they were able to turn off about a third of the hoods before they even had to consider any efficiency measures. So finding an advocate is very key.

The other thing is, again, just good old-fashioned awareness. It is amazing. When people have actually surveyed users on efficiency opportunities and how you can manage the hood, there are seasoned veteran users of hoods who thought – who have been confused about when to actually close the sash or not. They thought that "Oh, really? I thought that it is safer when it is open," when in fact it is actually safer when it is closed. So just awareness training is good.

The other is you can do things like simple stickers on hoods that tell you how much energy you are wasting, for example, when you leave a sash open. Or if you want to use the carrot, you can say how much energy you are saving by closing the hood. So, again, not a lot of rocket science here, but the classic measures of finding advocates and doing some awareness, as well, are ways to engage.

I will just mention one other area, because I mentioned hoods. I think another area is that – another area that scientists can really play a big role is in the procurement side. Again, plug loads and process loads in labs are a big chunk of your total energy use. Can be up to 50%. I mean, most office buildings, it's only about 25%. So you do not want to treat that as just something, oh, we have no control over as facility managers or as designers. We want to really take that bull by the horns.

And so there, I think you can work with scientists and say, "Hey, if you're ordering a piece of equipment, are there functionally equivalent options, alternatives, that may in fact consume less energy?" And a good place to start, frankly, are ultra-lows. Labs21 has developed several guidelines, and we also have a lot of data on ultra-low freezers. And, frankly, there are different options there in terms of how efficient or inefficient they are. So I think that would kind of summarize two areas that you can get users engaged.

Kathy Hyland: Okay. Next question is to Will. Will, the complexity of research facilities is daunting when it comes to trying to think about meeting your energy savings goals. Are there methods that we can use to break down activities into smaller, more manageable pieces?

Will Lintner: Well, we touched upon that in the seminar here. We would start off with a progression, starting with benchmarking, then going to our profiling tools and getting the potential ECMs. The next step would be to go to the more detailed assessments, and these would be done in connection with some expert in the field.

In the final analysis, when we do basically the investment for the ECMs, we would most likely want to go to the firms that specialize in this type of work, and there are some firms that do. And for instance, our case studies sometimes list those firms, and the websites that we have will have the case studies, and they also have our certified energy practitioners listed. And that is a good resource to go to the final level, where we have investment decisions to make.

Kathy Hyland: Thank you, Will. We have another question from Syracuse, New York, on the line.

Audience Member: Yeah, the question is regarding labs. When it comes to ventilation systems in labs, have you run into any issues with trying to comply with OSHA guidelines on PELs or industrial hygiene guidelines, in terms of controlling the ventilation?

Dale Sartor: Well, that is a good question. It is interesting, there is actually not a whole lot of requirements or standards relative to ventilation in labs. There is a requirement for ventilation when hazardous in an H-occupancy, and that requirement is one CFM per square foot. And if you have a ten-foot ceiling in a laboratory, that is the equivalent of six air changes. There is also some requirements in the fire code relative to air changes. So, certainly, codes and standards are a player, but, again, there is usually many opportunities to work within those requirements to reduce the air change rates or the amount of ventilation.

Kathy Hyland: We have a caller from Detroit, Michigan, on the line.

Audience Member: Okay, I have question regarding the power distribution chain. During one of the presentations, a slide was shown showing the power – the unit of power going from the power plant through transmission, through distribution, to the actual microprocessor device. We talked about the 80 Plus efficiency improvement on the power supply within the server. So my question to you all is: What is the federal energy facility management team's position or general feedback on simplifying that power distribution chain by going DC power distribution? So you would eliminate the multiple stages of conversions and PDUs, and a rectifier inside a server – combine that into a simplified equipment by distributing high voltage or low voltage, depending on who you are speaking to, right, the audience, directly at a data center facility. Thank you.

Dale Sartor: Thank you for the question. LBL has been doing a fair amount of research in DC powering of data centers. The theory is there. I mean, if we can go directly from the UPS system to the motherboard without converting back and forth between AC and DC, there is power to be saved.

Some of the issues are in the details. We are trying to develop a standard voltage. It looks like that is going to be 400 volts nominal, maybe 360 to 400 actual. Another problem or challenge is a standard plug. Everybody wants to plug anybody's server into their rack. So we are working with industry to develop a standard for a plug. There have been full-scale demonstrations and pilot demonstrations of DC powering data centers at higher voltages.

So I think it is a good question, and it is something to watch for in the next few years. In fact, a number of major manufacturers, including IBM now, have introduced high-voltage DC servers. DC servers have been around for many, many years in the telecommunication industry, so this is not a new technology. But the voltage that telecoms used was 48, and at the server density that we have now, we really could not afford the cost of the conductors to run power at 48.

So DC powering is a viable technology. A higher voltage of DC is coming. It makes interface with renewables and on-site generation easier. So I think you are onto a great concept, and it is an emerging technology that we are close to being able to deploy.

Kathy Hyland: Okay. That concludes our seminar for today. Please take a moment to complete a brief evaluation to help us determine what future training topics you would like for FEMP to offer, and also ways we can improve these First Thursday Seminars. You can also complete a quiz to reinforce your learning, and then you are able to print a certificate for your records. You can access this quick evaluation and quiz in one of three ways. Go to the website, www.femp.energy.gov/firstthursday, and find the quiz and evaluation there. If you registered for this course, you will get an email, and that will contain a link to the evaluation and quiz. And if you are watching this today by live webcast, you can click on the paperclip icon, and it will take you to the evaluation and quiz.

Thanks to Dale, Paul, and Will, and we will see you next month, on Thursday, August the 4th.

[End of Audio]