Daniel Robinson: Steve Greenberg is a senior energy management engineer at Lawrence Berkeley National Laboratory. He has researched and applied energy-efficient building in industrial systems for a variety of clients on three continents over the past 38 years. His data center spans over 20 years and includes data centers ranging from server closets to supercomputer centers in excess of 15 megawatts.

He has been involved in design, design review, commissioning, and retrofit; and has taught workshops on data center energy efficiency for the Department of Energy, utility companies, and ASHRAE. He holds a bachelor's degree in mechanical engineering and a master of energy and resources, both from the University of California at Berkeley.

I'd like to thank Steve for joining us and thank you all for being with us today. Next slide, please.

The Better Buildings program has two partner programs, the Challenge and the Accelerator. The Challenge participants have pledged to reduce energy consumption in their portfolio by about 20 percent within ten years. Challenge partners here. Next slide, please.

The Accelerator participants have pledged to improve the energy efficiency of one or more data centers of 100 kilowatts or greater IT load by at least 25 percent within five years. So you can see that we're talking about – we can be talking about pretty small data centers with only 100 kilowatts.

So before we get started with our presentation, I want to remind our listeners that we'll take questions throughout the presentation via the chat function on the webinar application, and then we'll have a Q&A session at the end after the presentation. Please send questions through the chat box, and we'll try to get to them as we can. This session will be archived and posted to the web for your reference.

Steve, are you ready to help us with our small data centers?

Steve Greenberg: I'm ready.

Daniel Robinson: Great.

Steve Greenberg: Next slide, please. All right, welcome, everyone. This afternoon's session is going to first touch on why small data centers are important, and then we're gonna go through a number of opportunities for energy efficiency improvement in those centers, starting with simple measures in the information technology, or IT, area; in the cooling system area; and in the area of electrical distribution.

Then we'll get a little bit more complicated into the IT replacement activity and virtualization, and then we'll get into a higher level still of some capital investment requirement, but investments that are generally very cost-effective, including moving the functions elsewhere, doing power monitoring, and doing some fancier items to the cooling system.

And then we'll touch on the importance for training for IT and faculty staff, and review some resources that are available to help you in your energy efficiency efforts. Next slide.

All right, why are small data centers important? Well, okay, small is a relative thing. So it could be anything all the way down to a server closet and up to 5,000 square feet of computing center space. That's our cutoff.

And in those centers, they have nearly half of the total servers that are in use today, and they use something like 40 billion kilowatt hours per year in the United States. At $0.10 a kilowatt hour, that's $4 billion. A billion here and a billion there, as they say, and you're starting to talk about real money.

Embedded data centers are very common when we're talking about small data centers. And what's an embedded data center? It's one that's located within another building that has other functions, typically an office building. And even though they may be a small fraction of the space in that building, they can easily dominate the entire building's energy use because they're ten to a hundred times more energy-intensive than typical office building space.

So small data centers are also important because, as Daniel mentioned, they're a challenge. They're small, so they don't have a dedicated staff. It's bits and pieces. Different people are responsible for different things, and so it's no one's fulltime job.

There's security risks often associated with them because there isn't enough attention being paid to physical security for sure, but even more importantly to the information security. Typically the energy use isn't monitored in these spaces, so nobody really has a handle on what's going on.

That said, they typically have large opportunities. In the cases that we've looked at, they typically have a 20 to 40 percent overall savings opportunity, and a 30 percent savings, just to pick the midpoint, would result in something like $1 billion a year in savings just in the United States. Next slide, please.

All right, so let's get into the simplest measures. Those include turning off unused servers, improving the power management within the servers that remain operating, improving air management between the cooling equipment and the IT equipment, and then increasing the temperature set points toward the high end of the range set by the American Society of Heating, Refrigerating, and Air-Conditioning Engineers, more briefly referred to ASHRAE.

And then turning off active humidity control, which is consistent with the bullet above, and then minimizing requirements for uninterruptable power supplies, or UPSes. And then taking those one at a time, next slide, please.

Okay, unused servers. They're often known as comatose or zombie servers or just idle. These are servers that they exist in the data center; they use power, space, and cooling; but they do no useful work. There have been surveys, and it's estimated that between 20 and 30 percent of servers that are in data centers are such comatose servers.

And they use about 50 percent of the power that a fully utilized server would use and about 75 percent of the typical server's power, one that's being used at only 25 percent of its potential utilization.

So how do we attack this problem? We need to know what's running on each machine in order to know what machines don't have anything running on them. That inventory needs to be established and maintained, and then once we know which ones aren't being used, we can shut them down and remove them from the space, which frees up that space for other IT equipment. Next slide, please.

All right, improving server power management. There are two facts about server power management. One of them is that most servers are shipped with power management turned on, and most servers in use have had their power management turned off.

The resulting recommendation is fairly straightforward. It's check the power management settings and enable them. They exist in three places: in the processor, in the operating system, and in the BIOS. Your IT staff will be familiar with those. There's an illustration of an example of where the power management has been disabled in a server, and that's an all-too-common fact and one that's easily fixed.

There are IT folks who just as a matter of course disable the power management because at some point along the way something interfered with it or it interfered with something, some application they were trying to run, and so as a result there's a tremendous amount of excess energy use that's unnecessary. Next slide, please.

All right, improving air management. So what do we mean by that? Typically these data centers are air-cooled. Pretty much 100 percent of the equipment in them is air-cooled. In the illustration on the lower left showing servers in a rack, normally there would be many more servers in a rack, but just as a representative illustration.

They take cool air in one side, heat it up by virtue of the computing that they're doing, and blow that air out the hot side. That hot air circulates back to let's label this a CRAH here, computer room air handler. It could be a CRAC, which is a computer room air conditioner, or just a simple air handler. In any case, that's the cooling equipment that cools the air back down and circulates it back to the inlet side of the IT equipment.

So what ideally happens is that all the air coming from those CRAHs or CRACs goes to the inlet of the IT equipment, and all of the outlets there coming from the IT equipment circulates back to the cooling equipment, the CRAH or CRAC, with no mixing at all between the hot and cold air. But most data centers are fairly closely represented by this example where the airspace is just open. So the likelihood that there'll be no mixing is fairly small.

So you have two problems. In the lower right, the equipment rack on the left, you have the case of recirculation where you have more air flowing through the IT equipment than being supplied to it. So you have hot air recirculating from the hot aisle back into the cold aisle, and heating up the air that's being used for cooling. Obviously that's not a good thing.

You also have bypass where you have more cool air being supplied than is being drawn through the IT equipment, and so you get cold air bypassing the IT equipment and then mixing with the warm air on the way back to the cooling equipment. That's not as bad from an operational point of view of the IT equipment, but by diluting that air it reduces the capacity of the cooling equipment and requires excess fan energy relative to what you would normally need. Next slide, please.

Okay, so how can we improve these things? Well, again it's fairly straightforward conceptually. We want to clear the desired air path. For example, these two illustrations of the one on the left of an underfloor area. Again, the underfloor in raised-floor data centers, that's where the cool air is being supplied from.

If you've got it jam-packed with cables like this, you can see that the air is gonna have a hard time getting to where you want it to go, as opposed to the one on the right where excess cables have been removed. There's a cable-management program in place. The excess cables have been removed, and the cables that are left have been nicely bundled so that they provide a minimal resistance to airflow.

So in addition to opening up the desired air paths, we want to block the undesirable air paths. Within racks, blanking panels, the most mundane thing, these are less than $5.00 a pop. They're plastic. Anybody can pop 'em in there. They make a huge difference in terms of internal recirculation within the rack.

And also cable and conduit cutouts under the floor and into the ceiling plenum, if one is being used. Those are often places where air is flowing where you don't want it to. And then the rack tops and row ends, strip curtains or more formal rigid barriers are options there. Those are not really in the simple category anymore, but that's a more advanced form of air management. Those are options, too, in order to reduce the mixing. That's what good air management is all about.

And why bother? What are the benefits? You can raise the supplier temperature, which reduces the cooling energy, and still fully satisfy the cooling needs of the IT equipment. And you can also supply less air because there's less air bypassing the IT equipment, and so you need to supply less in the first place and that reduces fan energy. Next slide, please.

All right, so increasing air temperature set points for the high end of the ASHRAE range. Okay, what do we mean by that? Well, the only temperature that matters in the data center is the temperature of the air entering the IT equipment because the IT equipment then heats it up and spits it out the back and it gets cooled again. But the only thing that really matters to the IT equipment is the temperature of the air entering it.

So the recommended range of ASHRAE is from 65 to 80 degrees Fahrenheit. That's slightly rounded, but that's effectively the range. And it's important to note this is not the same as the temperature set point for the cooling equipment because the cooling equipment is generally controlled either by the air returning to it, which would be the warm air, or less of the time, it's controlled by the air that it's delivering directly out of its bottom and not the air entering the IT equipment.

And so if you've got a CRAC or CRAH control based on return air temperature, that temperature will be much higher than the temperature entering the IT equipment. So one needs to take that into account. But before changing these set points, one needs to have good air management or the likelihood of creating hotspots is pretty high.

So by doing this, by raising that temperature, you can raise the temperature of either the chilled water or the CRAC – the temperature at which the CRAC compressor is absorbing the heat. That saves energy because they're pumping that heat less uphill, as it were. There are thermodynamic savings which translate into compressor energy savings in your cooling equipment. Next slide, please.

All right, turn off active humidity control. Okay, so back in the day when I was a lad, computer centers had a lot of paper in them. They had punch cards and they had fanfold printouts, and paper is very humidity-sensitive.

Well, that's all long gone, and so ASHRAE has, as of 2015, significantly expanded the recommended and allowable ranges for humidity in data centers to where in most cases active humidity control can be turned off completely. So that recommended range is basically 16 to 59 degrees Fahrenheit dew point temperature, and "dew point" is a measure of absolute humidity, and also 60 percent relative humidity. Again, those are conditions at the IT inlet.

So the problem that we typically see is that the cooling systems, the CRACs or CRAHs, they're often set to a much tighter humidity control than that, and they're based on the return air, which is usually far from the supply air and again from the IT inlet air temperature.

So as a result, they're often dehumidifying – some units will be dehumidifying while others are humidifying. They'll be fighting each other, and that results in a lot of extra energy use that results in no benefit. And so greatly expanding the control range or turning it off entirely results in typically both humidification and dehumidification savings.

Just an anecdotal case study that was here at the Lawrence Berkeley National Laboratory. We disabled the explicit humidity control. The cooling used almost 30 percent less energy, and the humidity barely changed because some units were humidifying and others were dehumidifying. Next slide, please.

Okay, uninterruptable power supplies. In order to have a secure or to enhance reliability, uninterruptable power supplies are there to take over in case a utility fails. Maybe they're backed up by generators, and maybe they're not. That's another whole area that we won't get into in this presentation, but it's an attempt to make the power supply more reliable.

Unnecessity redundancy leads to inefficiency. So what does that mean? This graph provides an example of it. When you have a redundant operation – so many data centers have their IT equipment – each server has two power supplies. So they're double fed. They feed each of those from power that's on an uninterruptable power supply, and they share the load.

That means that at most you would have 50 percent of the load on each UPS. Normally if one feed fails, then everything switches over to the other, and it suddenly gets bigger. But the reason it's 40 percent is there's the 80 percent rule for loading of electrical equipment that's continuously loaded.

It's not allowed to be used at higher than 80 percent of its nominal capacity. So when you have a double-fed system, then the highest normal operating load is 40 percent, and typically much lower because the system is significantly oversized.

So things to remember about uninterruptable power supplies is that many applications can be shut down and restarted without adverse effects. They aren't so critical that they need to be available every second of the year. Critical applications, if there are such applications, can be – should be considered to be moved to a larger data center or the cloud.

And then beyond that, once the small data center – once those critical loads have been identified, one needs to carefully look at the UPS needs to minimize the number and size of them. There are Energy Star UPSes available now that are substantially more efficient than the legacy UPSes.

There's an eco-mode option which turns off some of the conversions in the UPS until they're actually needed and very quickly activates the battery backup. Typically they bypass the conversions to and from the battery, but in less than a cycle they can identify a problem with the power and turn on the inverter output from the batteries and save those conversion losses the rest of the time.

And so often savings, they vary tremendously depending on the details of the center, but it's not at all uncommon to find ten percent savings, and in extreme cases up to 50 percent, because you can see if you're loaded in the ten percent range, and this is not that uncommon, your efficiency might only be in the 60 percent range, which means that 40 percent of the power is being lost. So big savings are definitely possible in this area. Next slide, please.

All right, a little more work. Okay, refreshing the oldest IT equipment with new, high-efficiency equipment. Also consolidating and virtualizing applications. Next slide.

All right, refreshing IT equipment. New equipment is more powerful, doing more computing in an absolute sense, but also is more efficient in that it does more computing per watt than older equipment, and it has better power management features. So that extra computing power means there is more virtualization potential, and a thing to keep in mind here is that because of energy and software cost savings you can refresh faster than you otherwise would.

So how to go about that? In procurement, select Energy Star servers, networking equipment, and storage. Use of solid-state drives versus hard disks is an additional option that saves some energy.

And then 80 PLUS power supplies. So Energy Star requires a certain level of efficiency of the power supplies, but one can specify a higher level like gold or platinum or even titanium beyond the bronze that Energy Star requires. So there's an additional five to 20 percent savings possible depending on the efficiency of the power supplies you've got now. Next slide, please.

All right, so consolidating and virtualizing applications. As we mentioned before, most servers operate with very low utilization, and utilization is just simply how busy is the processor. If it's 100 percent loaded, that means it can't do any more computing.

Well, most servers are luffing along at five to 25 percent, five to 15 here. That's typical for most servers. As we mentioned before, at those utilization rates, they're using three-quarters of what they would if they were fully loaded.

So what that means is there's an opportunity for virtualization. What does that mean? It's running multiple software applications on one physical machine. So you could run several – what had been several servers' worth of applications can be run on one server, and so that obviously saves space. It saves power and cooling, and there's just less equipment to keep track of and maintain.

And so what that also facilitates is consolidation. Once the number of physical machines is reduced, then there's more space available, which means several small data centers could be consolidated into one, that sort of thing. Next slide.

Okay, bouncing up a level, a bit more difficult but still generally very cost-effective. Moving applications or hardware to higher-efficiency data centers or to the cloud, implementing IT and infrastructure power monitoring, putting variable-speed drives on cooling system fans, installing rack or row-level cooling, using an outside air economizer, and implementing dedicated room cooling. Next slide.

All right, so typically it's possible to save a significant amount of energy by moving applications or sometimes the servers themselves to a larger data center – that's the consolidation concept that I just mentioned – to a colocation center, or to the cloud – to the greater cloud.

Typically the benefits of that are better security, both physical and information; better redundancy because those data centers typically have better power and cooling resilience; and better efficiency because they're very – they have relatively good management because they're big and their energy bills are big and they're paying attention to it.

So in evaluating these options, one needs to consider mandates. For example, some facilities are required to be completely self-sufficient in the event of infrastructure failure. They've got uninterruptable power supplies and generators so that their data center can stay up even if there's no network and no utility operation. So hospitals and such. There are certain types of facilities that can't really exercise this option, but most can.

Another thing to consider is the cost of making the move and then the total ongoing cost of staying versus moving. So if you're farming out this service to the cloud or leasing space in a colo data center, all those costs need to be taken into account. We call it the total cost of ownership needs to be compared. Next slide, please.

All right, so implementing IT and infrastructure power monitoring. So how much energy does this save? Well, it doesn't save any by itself, but it lets you know how your data center's performing, and so it can help not only identify opportunities but also check how well you're doing on an ongoing basis. They say you can't manage it if you don't measure it, and I think there's a lot of truth to that.

So what does one want to do in a data center? Well, tracking the performance of the power and cooling systems, and monitoring the information technology loads. And how does one do that? By installing some meters in strategic locations or reading meters that are already there. For example, most UPS systems have an output meter. Maybe it's networked; maybe it's not. Those are also opportunities.

And so it lets you get a handle on metrics that can be useful for tracking performance. One of those is of course power usage effectiveness, which is the total data center energy divided by the IT input energy. So when one thinks about it, if 1 would be some hypothetical ideal, 1.0 – and over 2 is getting pretty bad – there are a lot of data centers in that range, and so there's a lot of opportunity out there.

Big opportunities if you're over 2. If you're under 1.5, you're good. That's the DCOI requirement for existing data centers. And under 1.2 is really good, and there are data centers out there that are achieving under 1.2. There are ones out there achieving under 1.1.

So there's a guide to help you with this process, and there's a link to it. Next slide, please.

Great. Okay, another opportunity is putting variable-speed drives on cooling system fans. So your computer room air conditioners or computer room air handlers, they typically have constant-speed fans, which means they supply constant airflow whether it's needed or not. As we mentioned before in the air-management section, airflows are typically substantially higher than they need to be, and that's especially true once you have good air management.

So because the fan power is not linear with respect to flow, a relatively small reduction in flow results in a relatively large reduction in fan energy, for example, 20 percent resulting in 50 percent savings. That's a very commonly possible scenario.

And then because that fan energy is energy that contributes heat to the data center that needs to be removed by the cooling system, there are ancillary savings. And so you get between 20 and 30, round numbers, savings in the overall cooling system energy, so both fan and compressor energy and heat rejection energy as well at your tower or your condenser that's outside. There is a case study that we did, and there's a link to that. Next slide, please. Thank you.

Okay, so another opportunity is to install rack and/or row-level cooling. This only applies if you're replacing your racks or they're newly installed racks. I guess it is possible to retrofit this also on existing racks, but that's harder.

So what does this do for you? Okay, it's moving the cooling closer to the source of heat, and the closer you get, the more efficiently the cooling system can remove the heat, and that results in energy savings. There are various types of such cooling.

There's the in-row style where you have like every third rack or so is a cooling unit that takes air out of the hot aisle, cools it, and blows it into the cold aisle. There's rear-door schemes, such as the one that's illustrated, where the servers blow air through that door, which is a thin _____ *[break in audio]* and-tube heat exchanger like a car radiator that cools the air back down to the inlet – back down to the IT inlet temperature right there at the rack.

And then you have in-rack systems where the air doesn't flow in and out of the rack. It just flows around within the rack, and there's a cooling coil provided in the rack either on the side or typically below the IT equipment. There's internal circulation within the rack itself, and so then there's no need – if your whole data center is equipped this way, no need for CRACs or CRAHs in that space because all the cooling is happening very locally to the IT equipment.

So what are the benefits of this? The closer the cooling is to the IT equipment, the more efficient it is, because the whole bypass and recirculation problem of airflow is solved because of this tight connection. So less excess heat is generated in the cooling process, so that's less heat that needs to be removed.

And then because this cooling is so efficient, the temperature of the water used for these – and they're basically all water-based, although there are refrigerant-based ones as well, but typically they're water. The cooling plan is more efficient, and you can – if you have a water-side economizer, you can get more hours of re-cooling. Next slide, please.

All right, an outside air economizer, also known as an air-side economizer. What does that do? It uses outside air when conditions are suitable to provide cooling. So a downside, especially for retrofit, is you need a lot of air for data centers, and so you need an outside wall or roof for adequate access to large airflow.

An air-side economizer can be done with a built-up air-handling unit, a CRAC or CRAH with outside air capability, or just an exhaust fan with an ability to allow inlet air as well. So you can have a CRAC or CRAH circulating air within the space, but you have a separate exhaust and inlet that can allow an economizer function to occur.

So what happens is, as the illustration shows, you have outside air coming in. It may be mixed with return air. Often the outside air is too cold. It's below the allowable range or below the recommended range, and so you mix in return air, the warmed return air from the IT equipment, to get the desired supply temperature. And then the relief fan is shown at the bottom for letting the air out to maintain proper pressurization.

So the result is that the compressor, whether it be the CRAC compressor or the chiller compressor in the chiller plant, will run a lot less, and there'll be large cooling energy savings as a result. Next slide.

All right, dedicated room cooling versus central building cooling. So a dedicated unit will let you turn off – this is especially applicable to embedded data centers. You've got your 5,000-square-foot data center in a 100,000-square-foot office building. But the entire office building plant needs to stay on in order to satisfy this data center.

Well, that's a real tail-wagging-the-dog situation. So you let the rest of the building operate on a normal building occupancy schedule, and you run the dedicated cooling system just for the data center.

So how to do it? Well, figure out if it's applicable, of course. And then in selecting it, pick a high-efficiency unit, one that has a high seasonal energy efficiency ratio and one that has an outside air economizer for reasons that we just went over. And then control that unit based on the IT inlet temperature because again that's the temperature in the data center that matters. Next question. Sorry, next slide. Thank you.

Okay, so training. Training is very important for people to get up to speed and stay up to speed on opportunities and best practices. There are various places to get such training. Utility companies can be one of them. Some companies do offer such training for their service area.

ASHRAE, again they have a training series. FEMP has a training series, including the one that you're in now, and the Center of Expertise for Energy Efficiency in Data Centers, there's training available there. And the Data Center Energy Practitioner program, as noted previously, is a requirement of DCOI, and there's a link to that. Next slide, please.

All right, so some resources that you should consider. There is a guide that's the companion to this webinar that we produced last year, and please avail yourselves of that. Next slide.

All right, there's a couple of centers, including the Better Buildings Solutions Center and, as mentioned before, the Center of Expertise, and links to those. Next slide.

Okay, and even more resources. On the Center of Expertise site, there's air management tools. There _____ *[break in audio from 0:41:49 to 0:41:59]* case study we mentioned. Those are all there. The Data Center Optimization Initiative, no doubt many of you are already aware of that. There's a link to that. And then of course Energy Star equipment, which we mentioned several times, as a resource for finding efficient IT equipment as well as UPSes. Next slide, please.

All right, so as noted here, the slides – as mentioned, the introduction slides will be posted online, and you can contact Daniel or myself with further questions. But we have some time now to take questions from anyone who's got one. So thanks very much for your attention, and let's continue the dialog.

Daniel Robinson: Great job, Steve. I really appreciated that walkthrough for small data centers, and we have a bunch of really good questions. I'm gonna give you a few that I think you can answer quickly to make sure everyone's on the same page. You referred to A1 through A4 for the ASHRAE standards. Could you discuss that briefly?

Steve Greenberg: Sure. So those are thermal guidelines for data centers. A1 is the most restrictive. One can specify your data center to be an A1 data center, and so vendors supplying IT equipment to it know that if their equipment operates satisfactorily in A1 that they'll be fine.

If one selects the higher levels, which are higher temperatures and broader humidities allowed, you can tell vendors, "I have an A3 data center, so your equipment needs to be compatible with that." And vendors will supply such equipment.

It might be a little bit more expensive, but there are big savings in construction and operating cost because you're running smaller cooling systems and ones that may not even have compressors, for example, because they can use some combination of outside air and water-side free cooling to meet those loads.

Those are those categories, and in that thermal guidelines document that's shown on that slide, the ASHRAE thermal guidelines document, it gets into great detail on that topic.

Daniel Robinson: Great. Thank you, Steve. I'll try to put a couple of questions here together. Power management. Do you find that IT managers are resistant to taking UPSes offline, and if so, what's a good way to convince them? Particularly data security or resiliency of the data center is gonna prevail over energy efficiency objectives typically.

Steve Greenberg: Right. Yeah, that's a very valid concern, of course, that the UPS is there to provide redundant power. It's kind of a relative thing. Yes, most operators have some reluctance to operating their UPSes in the so-called eco-mode or bypass mode because they fear that it won't successfully recover.

If you've got a double-fed data center, then one can operate one of the UPSes that way and not the other one, for example. There's intermediate schemes that sort of mitigate that risk. Everyone has their level of acceptable risk, right?

Another trend that's happening, besides moving to the cloud, is having the redundancy not in an individual data center, but across data centers, because it's recognized that any one location may be taken offline by extreme events, if the computing functionality is really that critical that it needs to be backed up somewhere else. And then your redundancy is provided in that way rather than sort of piling on extra sets of belts and suspenders at one particular location.

Daniel Robinson: Great. Another power management question. Do you have any estimates for savings from turning on server power management? And another similar question. If the power management settings are turned on in the BIOS and the OS and the processor, will there be any savings or is it an all-or-nothing situation?

Steve Greenberg: Answering the second question first, it's definitely incremental, and there are also levels. In some cases it's not just on or off, but it's like how long does it take for an unused disk to – an idle disk to be turned off, for example. That requires recovery time to spin the disk back up if the information on it is needed. So there are various intermediate settings, and then whether you do any or all gives you quite a range of potential savings.

In getting to your first question, how much can you save, there's a huge range, but in the tens of percent are possible, probably the low tens. But again, it depends on how utilized the IT equipment is to begin with. The power management reduces the power to it when it's operating under a lighter load or no load, and so it's a huge function of how heavily loaded it is in a –

Daniel Robinson: Can you provide – oh, sorry.

Steve Greenberg: – in a compute utilization – from a compute utilization point of view.

Daniel Robinson: Great. Can you provide us with an idea of the typical first cost for various recommendations as well as a typical payback?

Steve Greenberg: *[Laughs]* There were quite a number of potential measures covered, and so the short answer to that is not in the next six minutes. But there are quite a number of case studies that are covered on the center website. But every case is different.

One of the reasons we selected these is that they are typically cost-effective in small data centers because the opportunities are so big to where a lot of times a really small data center it's not worth doing the analysis to figure out whether it's cost-effective each and every time. But if one has a number of such centers, it's worth just implementing the measures because on average, they'll be very cost-effective.

Daniel Robinson: Okay, great. Regarding small data center energy management, is there a particular technology – is the technology for large data centers scalable to small data centers?

Steve Greenberg: I would say it varies. Quite a number of these opportunities apply to any data center. So the free cooling, the close-coupled cooling, doing good air management, all those scale everywhere across the entire range of data centers from server closets to multi-megawatt facilities.

That said, there are certain ones where it's not really worth it. For example, the so-called data center infrastructure management packages that a number of vendors have available these days, it wouldn't be worth trying to implement one of these full-blown ones to a small data center because the overhead costs of getting it set up would greatly outstrip the benefit in general.

That's for the full package. But there's still – for example, keeping track of applications running on each server, that's worth doing no matter how small your data center. That's a form of energy management. Same with monitoring the power management features in the IT. That's completely scalable as well.

Daniel Robinson: All right, so there are a number of questions about how – what kind of tools are available to estimate savings. I thought I'd tee that up for you to point out some of the tools that you mentioned earlier.

Steve Greenberg: Right. The data center profiling tools, there's a couple of those. There's the full-blown DC Pro, as we call it. That's available on the data center – the Center of Expertise for Energy Efficiency website that we put in there.

There's also a quick-and-dirty PUE Estimator version of that that will give you a quick PUE estimate. If you don't have the metering installed, you can at least get a good guess as to where you're starting from.

In the case of the full-blown DC profiling tool, there's – it gives you not only estimated current PUE, but how much better your PUE could be if you implemented certain measures. It gives you an estimated breakdown of usage as well.

I implore people to go to both of those sites, the Better Buildings site and the Center of Expertise site, because there's a lot of valuable resources available at both for free.

Daniel Robinson: All right. I'll try to take advantage of the next couple minutes to deal with a few questions here – quick questions here with cooling. What other air-based cooling technologies can be of value, in particular liquid cooling?

Steve Greenberg: What other air-based cooling technologies?

Daniel Robinson: Sorry, other than air-based.

Steve Greenberg: Oh, other than air-based. Right.

Daniel Robinson: And then also on that note, what is the difference between in-row and in-rack cooling?

Steve Greenberg: Right. So all of those close-coupled coolings we talked about use liquid close to the IT equipment, but one can actually get off-the-shelf now and you can also convert existing IT equipment to use water cooling to the chip, which is even more direct, even closercoupled, and results in even greater savings.

There are also immersion cooling technologies for the real adventuresome folk where you dunk the IT equipment into a bath of non-conducting fluid, non-conducting liquid. That's yet another opportunity rather beyond the scope of this webinar.

Oh, the difference between in-row and in-rack. They're very similar except the in-rack coolers have the cooling coil – the chilled water cooling coil within the rack where the IT equipment is.

And in-row, it'll be another cabinet next to the IT equipment rack, but right next to it in the same row. So it circulates air. It still uses the hot aisle and cold aisle, but it's very local circulation rather than the bigger circulation in an open data center using CRACs or CRAHs. Clear as mud?

Daniel Robinson: I thought that was good. So one recent question. Most of these solutions have been available for years, but small data centers are a challenge to achieve a respectable return on investment. What kind of financial incentives and aid can you point us to to help overcome this obstacle?

I can point one out. On the Better Buildings blog, we posted about a Maryland state incentive for small data center new builds or retrofits. The state of Maryland would fund up to 50 percent of some small projects.

Steve Greenberg: Yeah, that's a great example. Other times utility companies will have rebate programs. Typically data centers fall into a custom rebate where you show savings and they'll give you a rebate based on those estimated savings. So check with your utility company for share. Yeah, state energy offices and utility companies are good places to check.

Daniel Robinson: All right. Well, Steve, thank you very much for guiding us through that presentation and answering our questions. If you could pass to the next slide, I want to point out that the Better Buildings webinar series will continue next month. Sorry, the previous slide showed some additional resources which you'll be able to click on the links in the coming week.

But the next webinar is Buildings that Rebound: Resiliency Strategies for Commercial Buildings and Communities, and I'm sure it will be great. Please join us for that, and thank you for participating and asking so many great questions. Have a great day.

[End of Audio]