E3T Energy Efficiency Emerging Technologies

Washington State University Extension Energy Program

Bonneville Power Administration

EnergyServices Western Area Power Administration

EMERGING TECHNOLOGIES SHOWCASE WEBINAR:

IT TAG DATA CENTERS: HVAC

Alan Matzka, Bradford Consulting Engineers

November 26, 2013
Question and Answer session

Q: What does CRAC stand for?

A: CRAC stands for computer room air conditioning.

Q: When utilizing improved equipment control to better satisfy real vs. design loads, what metric would you use to control the HVAC equipment. Would indoor temperature be one of the metrics?

A: Yes, I would recommend using indoor temperature as one of the metrics. For example, say you're buying equipment to put into a rack. Perhaps you're buying a Dell server or you've got legacy equipment. The temperature in the operating environment usually depends on the warranty of the equipment. If you have a problem you want to make sure that your warranty is still valid. Another factor is the stability of the operating environment for the amount of load you are putting on your server. Some legacy equipment can be run up to warmer temperatures. We all have this experience with the electronics we use in our houses – a computer or game console stuffed into a cabinet may get pretty hot but it doesn't necessarily break, although sometimes it does. If temperatures increase, the limitation is moved from the data center environment to the air conditioning equipment. Cooling to 65 degree air temperature vs. 75 air temperature vs. 85 air temperature will, in each case, present a different return temperature back to the air conditioning equipment. Some air conditioning equipment can handle warmer temperature better than others. A real robust system, for example a chilled water cooling system, generally has a little more flexibility in what temperatures can be run to cool. Unitary CRAC equipment, which is basically an AC unit with a little condenser sitting outside or on the roof running a straight refrigerant to the CRAC unit itself, is more limiting on the types of temperatures that can be run through it. It's a difficult question to answer – in general if you can run your data centers warmer it makes things easier. A robust control system that can manage the economization, the CRAC unit staging, which units are running against, and tie that all back to load monitoring so you know what loads you have, can help you optimize and make for much better energy consumption inside your data center.

Q: Most door coolers are used to address hot spots *in addition* to existing air cooling - so where is the energy savings?

A: The energy savings in the case of a door-based cooler can be used for supplementation, very correct. However, if you're designing or retrofitting a data center, you can, for example, run chilled water out to your chilled doors or in-row coolers. We had a project that we worked on where there was not an ability to do straight outdoor air economization, so we ran the cooling volume direct to the racks. One of the advantages is that the air circulation path is shorter, so the energy that's normally expended to drive air over bigger distances or through plenums is removed. So that's an energy savings. The cooling method implies chilled water, and it takes less energy to pump chilled water over the distance than it does to drive air down ducts. The energy savings, again, is going to be dependent on where and how you are implementing it. It depends on the systems you have deployed and if you're going to use it to supplement or if you're going to use it for wholesale conditioning of your data center. I hope that makes sense; it's kind of my take on it.

Q: But wouldn't it use even less energy if it avoided high power / high heat racks and avoided the hot spots in the first place - and thus avoid the need for spot coolers?

A: Yes, that would be true if you could even out the density. That would certainly be the case. So, if the choice was between doing average density that is lower versus having a lower average density for most of the racks and having some high density locations, then yes that would be better. We have to think about the practicality of how you apply loads to your data center. This is a whole separate subject matter altogether – trying to figure out how to tie provisioning of equipment in the data center, and the needs of the data side of the operation against the cooling and maintenance of the facility operations side of things. It can be a very complex question. I think that's one of the key takeaways from any technology discussion such as this – there are plenty of opportunities to find exceptions to the rule and make it work differently. I think the best way to approach all of these technologies is to look at them as an option; maybe not the best option, maybe not even a good option, but still an option worth looking at and an option that could potentially help us.

Q: Are there are any heat recovery strategies that companies use to save additional energy?

A: Yeah that's a great question. It comes back to a density question. The heat that typically comes off a lower density data center is what I call very low grade heat. Of the things that could make use of the heat that the data center produces, most domestic hot water applications are low enough in temperature that they might be able to get some benefit from it. The unfortunate thing is that a lot of data centers are strapped to buildings that don't have a lot of domestic hot water need. If you're operating a data center inside of a facility that has a broader range of occupancies, for example a data center floor that's part of a multi-story office building that has a fairly large need for other types of water for either domestic hot water or preconditioning for heat recovery, then you can certainly avail it and then it will make sense. Then there are data centers with even higher density where you are getting the temperatures up off the backside of the rack, such as Facebook over in Prineville. There they are running their data center so hot that their return temperatures are up around 114 degrees or more and you're almost to the point where you can put coils in there and run hot water off them for bathrooms – it's that warm. But that's not practical for everybody's application. It kind of depends on your building situation. I'm going to start to sound like a broken record here with it just depends, it just depends, but it's really true.

Q: In order to provide the ever increasing IT loads to end-users (web, email, online streaming, etc.), has the industry started to accept that they will need to look at liquid cooling to take advantage of better cooling ability over air?

A: Well you know what's funny, I had a discussion with a friend of mine about 15 years ago and he turned out to be a hell of a lot smarter than I was at the time. He was asking me if there was any way to build a water-cooling method for personal computers because he was a big video game player and he wanted to get the fans out because they were too noisy and he didn't like them, etc. At that time I didn't know how you could do it and I didn't think it would ever be worth it. The funny thing is that in the personal computing space it's now very common. There are loads of solutions for water-cooled direct equipment. But at the data center level, there aren't many solutions. I think the biggest barrier is infrastructure. There are some parallels with buying an electric car – you want to be able to go and plug your car in but there aren't that many places to plug it in. It's a chicken and egg thing. If you are creating a purpose-built data center from the ground up and specing out your own servers which are being built to your specific needs, and you want to do something a little bit more exotic, then I think you could do that. The other thing is that temperatures are pushing up so far in the environment temperatures and we've got equipment that is being built to run at 100 degrees in the data center environment and you still need to maintain warranty and reliability. At that point you can just blow air into the building and you are basically cooling it no matter where you are at, even Arizona. Trying to deploy a water solution out to every individual piece of equipment at that point is not really practical. There is a tip-over point where if you can you can run data centers up hot enough that just blow air in there and you are done. If you have really low temperature and low density then the existing room-based solutions are just fine. If you are somewhere in the middle then it can work if you can spec a whole new facility and really develop it. It just requires a lot of commitment from different parts of the industry, from both the server manufacturer level and the facility level to be able to facilitate something like that.

Q: Can you give an example of how to understand actual loads vs. nameplate loads.

A: The only way I know of to understand this is to get into your power distribution and provide monitoring so that you can look at what the actual draw is into your racks. There are different solutions to this, some that operate at the rack level such as smart plug strips, and things like that that allow you to monitor the power. The other thing you can do is some statistical sampling. First understand the usage, how it's being accessed and how frequently it's being accessed, what type of services you are running, are you running compute versus running data serving versus running hosted web access, etc. Go out and rather than install an entire monitoring system, do some sampling. Another way to manage it is to look at the temperatures into and out of your racks. You can compare that to just some overall power draw. So there are different ways to get that kind of granularity. If you have the wherewithal to put in a smart monitoring system to understand what the real draw is to different pieces of equipment, that can be very advantageous but it can also be very cost prohibitive so it really depends.