A Comparison of Room-, Row-, and Rack-Based Data Center Cooling Products

A Dell Technical White Paper

Dell │ Data Center Infrastructure

David L. Moss

August 2010

Contents

Tables

Introduction

There are a wide variety of data center cooling products on the market to choose from, but what criteria do you use when picking a data center cooling coil strategy? Ten years ago the choice was very limited. It was either a Computer Room Air Conditioner (CRAC) with much of the refrigeration cycle built into the unit, or it was a Computer Room Air Hander (CRAH)—a simpler chilled water version that is typical of larger buildings with a central chilled water plant. You had a choice of unit size (tonnage) and whether the air flowed up or down through the unit. Other than choices such as humidification or temperature sensor location, there weren't many more options from which to choose.

Now you are faced with several row coil configurations: over, on top, and between racks; self-contained racks with dedicated coils on the side; coils hung on the rear of the rack (both active and passive); and you even have fluid choices of water versus refrigerant. Each coil vendor has his unique way of expressing the value of his product's form factor. Efficiency is an important factor, but claims are often cloudy and rarely present an apples-to-apples comparison. It is often hard to ascertain the true efficiency differences with the way the products are presented.

- "Requires 15% to 50% less chiller capacity"
- "Reduce energy usage on cooling by 30%"
- "Can provide up to 90% less energy consumption for the primary cooling loop"

These are example claims that sound great, but they are hard to put in context relative to other energy expenditures in the data center. Just how much difference is there between room, row, and rack cooling? Is cooling efficiency your main criteria? Or is it cost? Are there other considerations you should be looking at? This paper attempts a fair comparison of efficiency and suggests other considerations which may be of equal or greater importance.

Efficiency

Efficiency can be evaluated by studying the energy expended in 3 different areas: the chilling process, air movement, and humidity adjustment. Product differences are most easily studied in chilled water facilities since the chilling process is accomplished remotely in a central location. The energy to chill can be easily separated from the energy to move air and the energy to humidify. With most coil systems, air movement is a necessary part of the solution, but humidity adjustments are wasted. If you can reduce or eliminate dehumidification (coil condensation) and humidification, why not do it? Some of the row and rack solutions avoid condensation altogether when they use a fluid distribution unit that automatically adjusts fluid temperature above the room's dew point. This puts these systems at an advantage over other systems running directly on cold chilled water. But if you can raise the chilled water temperature, you also gain this advantage with any of the coil systems.

You can improve the efficiency of any coil system by increasing water temperature, supply air temperature, and/or increasing return air temperature. Not only will this decrease chiller energy and increase the capacity at each coil, but it will also reduce the coil condensation penalty. The chilled water temperature setting is a conscious decision; an upward adjustment may be limited in your facility if it is a mixed-use building where the chiller also serves office space. Supply and return air temperature increases come easier in the presence of a containment solution. This statement is valid for any coil product—room, row, or rack. Once containment is in place and hot spots have been reduced or eliminated, supply air temperatures can be increased. The separation resulting from containment should reduce the mixing of hot and cold air at the rack level, but to increase return air temperature, a lowering of coil flow rate is needed. Without variable fan capability and the right volumetric control system, this would have to be done by turning off cooling units.

If you've taken steps to increase data center temperature, you've removed much of the potential for humidity to affect your energy expenditure. If the air temperature increase was also accompanied by an increase in the chilled water temperature, you affected the largest energy spend, which is the chilling process. The remaining difference between room, row, and rack coils is the energy spent to deliver air. Just how much difference is there in the fan energy of different cooling coil systems? Some would present the difference as significant. This excerpt from a Green Grid's paper suggests large differences between cooling coil systems:

> "Locating cooling closer to IT equipment can reduce data center cooling costs by more than 30% compared with historical approaches to cooling. Air conditioning unit fans are known to consume a significant portion of energy in most data center cooling systems. Mounting the cooling modules as close as possible to the source of heat—such as placing them directly above, alongside, or within high-density racks—reduces the distance that the fans must move air. This can provide up to 70% savings of the energy required to move the air."[1]

I happened to coauthor this paper, and I now think it is a bit of an overstatement if you really consider an apples-to-apples comparison. Instead, I like to look at a metric that compares how many watts it takes to produce a specific flow rate, watts/CFM (cubic feet per minute). A quick perusal through technical documents for most coil systems will produce both the wattage and the CFM data.

Table 1. Watt/CFM Values for Common Cooling Coils

	watts	CFM	watt/CFM	Source
Liebert FH600C	11000	17100	0.64	http://www.liebert.com/common/ViewDocument.aspx?id=1151
Liebert CW106	9000	17100	0.53	http://www.liebert.com/common/ViewDocument.aspx?id=1151
APC® ARC100	920	2900	0.32	http://www.apcmedia.com/salestools/MLAN-78NKNU_R1_EN.pdf
Liebert XDV + proportion of XDP	242	1000	0.24	http://www.liebert.com/common/ViewDocument.aspx?id=188
Rear Door Coil- XDR (server penalty plus XDP)	3	26	0.12	http://www.liebert.com/common/ViewDocument.aspx?id=188

These five products were chosen to show contrast. The first two are large air handlers that are essentially the same unit, except the newer version (CW106) has more efficient EC plug fans. Like the larger air handlers, the APC row cooler is typically fed directly by building chilled water. This is a contrast to the remainder of the solutions that are fed by an intermediate pumping unit. The XD products are refrigerant-fed, requiring a pumping unit for the refrigerant. Unlike the first three units, the pumping unit keeps the operating fluid above the dew point and removes condensation potential everywhere except within the pumping unit. The XDV is a top of rack unit, and the XDR is a passive rear door with no fans. A straightforward calculation may be applied to the first three products. Vendor documents list specific fan power at maximum flow rate, so "watts/CFM" is a simple division of the two numbers. The power number for the Liebert XDV was calculated assuming 190 watts for fan power plus 1/16th of the 832 watts used by the XDP pumping unit. The value for the Rear Door Coil is presented on a per server basis and is based on two things: (1) an increase in server fan power (1 watt

[1] "Seven Strategies to Improve Data Center Cooling Efficiency." *The Green Grid.* October 2008. http://www.thegreengrid.org/~/media/WhitePapers/White%20Paper%2011%20-%20Seven%20Strategies%20to%20Cooling_092809.ashx?lang=en

in this case) due to the added pressure of the rear door coil, and (2) a proportional amount of the XDP pumping power.[2]

The watt/CFM metric in Table 1 varies from 0.12 for a passive rear door coil to 0.64 for the older model down-flow CRAH. This appears to be a significant difference. It is helpful, however, to show what these numbers mean relative to a single server. Per Dell's power estimating tool, ESSA[3], a hefty configuration for a 1U Dell™ PowerEdge™ R610 consumes 406 watts and 26 CFM. As an interesting side-note, the internal fan power producing the 26 CFM is about 8 watts (or 2%). A simple multiplication of the watts/CFM numbers of table 1 by 26 CFM yields the amount of coil energy used to produce the volumetric complement for this representative server. The power ranges from 3 watts for the door coil to 17 watts for the older CRAH unit and is represented in Table 2 by dividing it by the server power of 406 watts. Represented in this fashion, it effectively becomes a server overhead number that can also be inverted to be shown in terms of efficiency by subtracting it from 1.

Table 2. Cooling Coil Overhead / Efficiency

	Overhead	Efficiency
Liebert FH600C	4.2%	95.8%
Liebert CW106	3.4%	96.6%
APC ARC100	2.0%	98.0%
Liebert XDV + proportion of XDP	1.6%	98.4%
Rear Door Coil- XDR (server penalty plus XDP)	0.8%	99.2%

This table brings the differences in coil-related fan power a bit more into perspective. If looking at the exact airflow complement needed per server, the server energy requirement is only augmented 1%-4% for the broad range of products presented. Even with an inefficient backend operation (chiller pumps, tower), the difference at the facility would only vary the same 1-4%. Presented in this manner, the overhead number is essentially a partial PUE. According to the Green Grid,

PUE= 1/DCiE= Cooling Load Factor (CLF) + Power Load Factor (PLF) + 1.0 [4]

The cooling load factor is defined as the power consumed by all the facility cooling systems divided by the IT load. The overhead numbers of Table 2 are partial contributors to the CLF and would show up as values ranging from 0.008 to 0.042. So, in an optimized scenario where facility supply closely tracks IT consumption, the difference in fan power for the products considered in Tables 1 and 2 only contribute variance in PUE of 0.034 (0.042 – 0.008). We all know, however, that very few coil solutions provide the exact amount of airflow and no more or no less. It is not uncommon for data centers to be grossly over-provisioned in delivered air volume relative to IT consumption. If, however, the data center is

[2] For further detail concerning server fan power increases associated with external resistance, see "IT Equipment Response to External Pressure," by David L. Moss and Robert B. Curti. 2009. http://content.dell.com/us/en/corp/d/business~solutions~whitepapers~en/Documents~server-response-to-pressure-12182009.pdf.aspx

[3] www.dell.com/calc

[4] "Green Grid Data Center Power Efficiency Metrics: PUE and DCIE." *The Green Grid*. 2008. http://www.thegreengrid.com/~/media/WhitePapers/White_Paper_6_-_PUE_and_DCiE_Eff_Metrics_30_December_2008.ashx?lang=en

delivering 100%, 200%, or more air relative to IT consumption, this fan overhead obviously becomes more significant. One argument for row and rack solutions is that because of their proximity to the IT load, their volumetric balance can be more closely matched. Containment solutions also provide the opportunity to strike a closer volumetric balance enabling the facility to lower flow rates through VFD adjustments or through simply turning off unnecessary CRAH units.

Applying this argument to a different platform, a 2U server typically uses more airflow than a 1U. For the same configuration, a PowerEdge R710 consumes 43 CFM. Because of the higher flow rate, comparable overhead numbers would increase by a factor of their flow rates (43/26). A table for the 2U comparable to Table 2 would have an overhead range from 1% to 6.8%. A blade system would fall somewhere between the overhead compliment for the 1U and 2U. Storage products typically have much higher flow rates; their overhead numbers would fall above the 2U numbers.

Unless airflow is grossly over-provisioned, the difference from the best to worst coil is only a few percent. We'll leave equipment capital cost discussions to vendors of those coil products, but we should point out that traditional CRAH units are typically less expensive than newer rack or row coil products. The raised floor, CRAC/CRAH solution is often discounted as "old" technology not capable of supporting higher densities and being much less efficient. That is simply not true if comparing optimized solutions (contained, higher temperature, optimized airflow). This brings the discussion to one of control.

Simplicity and Control

Unless you choose the rear door coil, the challenge for any other coil is to match coil air delivery with IT consumption. You might be tempted simply to design for intentional over-provisioning. After all, the most important aspect of a cooling system is to ensure there is no under-cooling. Why not over-provision to be safe? Don't forget, over-provisioning means that cold air is bypassing the IT equipment and at some point is mixing with hot return air. This mixing lowers the return temperature and ultimately lowers the capacity of each coil. If you are trying to optimize your coil unit capital expenses, you want to return warmer air to them. What are some of the ways to lessen the over-provisioning but still insure adequate delivery? The IT systems themselves have independent fan control. Most of the other coil products do have the ability to adjust flow rates, but how do you make these two independent control systems work in tandem? How do you optimize?

APC has a good control strategy associated with their row cooling units, offering a coil delta-T setting where they vary the coil's flow rate to achieve a specific drop in temperature. If you know the aggregate delta-T (temperature increase) through the IT equipment, then it can be matched effectively by the APC coil. If it is matched, the flow rates are matched. However, IT delta-T is not necessarily constant, and it is often hard to ascertain what the aggregate delta-T is in an aisle of IT equipment. Liebert also has varied options for their new row cooling product. One interesting control strategy is Liebert's SmartAisle™ solution. It can be implemented with row or CRAH/CRAC units operating with cold aisle containment. Imagine that the cold aisle is reasonably sealed off. A small hole in the containment would offer an opportunity for air to leak in or out. In a cold aisle scenario, the temperature should always be colder inside than outside of the aisle. A simple temperature sensor measures the leak. If it is warmer than the cold aisle, air must be leaking in. The coil units would be sent an instruction to increase flow rates within the aisle. A slight drawback with this system is that in most cases 100% of a CRAH/CRAC unit's air does not go to a single captured aisle. An adjustment made for one captured aisle may affect another. This method of control is somewhat similar to what Dell is doing with a future rack level containment product.

Dell's rack level containment establishes a tight duct to the raised floor and completely captures dedicated openings under the rack (either vent tiles, or holes cut per Dell specifications). The IT systems themselves serve to regulate flow through the rack and ultimately what comes out of the floor. Rather than the pressure under the floor determining how much air flows through a floor opening, the passive rack, driven by IT air consumption, drives how much air is extracted from the floor. It also ultimately affects the pressure under the floor. If the server's needs exceed what would normally have been driven out of the floor, the rack consumes more. In this case, this would serve to lower the pressure under the floor. Similarly, the rack naturally constricts any overage that might normally have come out of the floor. Since the rack has the ability to affect pressure under the floor, it can be paired with cooling units that are easily converted to adjust their flow rate based on raised floor pressure.

Driven by a requirement to establish a specific pressure, all cooling units work in tandem and run at the same speed to drive to the desired under-floor pressure. Imagine how simple this solution is:

- Racks take only what they need out of the floor
- CRAH units respond to pressure changes affected by rack consumption
- Variable fan speed throttles CRAH air delivery down to closely match the rack consumption and will also do this if the racks are dynamic in their consumption

It is easy to understand when a new CRAH unit is needed—with all CRAH units operating at the same speed, a simple monitoring of speed would indicate when an (N+1)-1 amount of flow rate is being approached (i.e., 10 CRAH units in the room, need an 11^{th} when VFD approaches 90%)

What about mixing containment products? Certainly, hot side containment is difficult to mix with cold side containment. With hot side containment the IT is open to room temperature. Coexistence with a cold side strategy would put the exposed inlets to risk. Different hot side products can typically coexist with each other as well as different cold side products. The Liebert SmartAisle product would work well with the Dell rack containment in a raised floor environment. If mixing these two products, you would default to the SmartAisle's control scheme letting the CRAH/CRAC units still control the aisle flow rates. The Dell racks would simply extract the amount of air they need, using up some of the available air being distributed to the closest contained aisles.

One other solution that can match air handler air volume with IT consumption is an active chimney containment system from Opengate Data Systems. The chimney contains pluggable fan modules capable of providing flow rates comparable to up to about 30 kW of IT equipment. The system measures pressure in the rear of the rack. By doing this, it can adjust the chimney fan speed to match the aggregate flow of the equipment in the rack. This is helpful in evacuating the air out of the rack, but there is an added advantage. By matching the IT flow rate, the chimney system now understands the rack's consumption. This information can be aggregated by the chimney systems and sent to the control systems of the air handlers. Understanding the total airflow needs from the racks, the air handlers simply adjust their aggregate flow rate to complement. At this time, Opengate has coordinated this feature with a couple of cooling vendors, but not all. This system is fairly expensive when installed per rack. It can, however, be used effectively for more than one rack. It may still prove cost effective since it is paired with CRAH/CRAC units that are cost advantaged over row and rack coils. In fact, any of the intelligent containment options that operate with CRAH/CRAC units (SmartAisle and Dell rack) should come out favorably in a cap cost comparison.

Cost, efficiency, and simplicity are the three metrics that make the most sense to this writer. If efficiency is important to you, you are likely to consider containment and the resulting energy-saving options like higher data center temperatures and better control of volumetric delivery. As for simplicity, the questions you should consider are whether or not the solution:

- Meets the minimum requirement in that it provides enough air to all IT equipment?
- Allows you to put equipment anywhere without fear of a hot spot or inadequate air delivery?
- Adjusts when adding new racks?
- Adjusts to dynamic IT operation, increasing air provisioning when rack air consumption increases?
- Helps you with coil capacity planning?

Variable Air Moving Capability

If variable speed operation is an option on the coil you are considering, choose it! It will give you control options like those mentioned earlier. It will help in capacity planning. For instance, if your IT systems are all receiving the right temperature and you know your coils are operating at, say 80%, you know you have 20% in reserve to cool more equipment. Variable fans will also save more energy. Relative to flow rate, fan power varies cubically. If you run fans at 90%, the energy required is (90% x 90% x 90%) 73%. A 10% reduction in flow comes with a drop of almost 30% in fan power. Let's look at what happens to the watts/CFM metric of Table 1 if fans are run other than at maximum speed.

Table 3. Watt/CFM Metric With Variable Flow

	watt/CFM →	90%	80%	60%	40%
Liebert FH600C	0.64	0.52	0.41	0.23	0.10
Liebert CW106	0.53	0.43	0.34	0.19	0.08
APC ARC100	0.32	0.26	0.20	0.11	0.05
Liebert XDV + proportion of XDP[5]	0.24	-	-	-	-
Rear Door Coil- XDR (server penalty plus XDP)	0.12	N/A	N/A	N/A	N/A

Since fan power varies cubically, the watt/CFM varies as the square because of the extra CFM term in the denominator. The overhead numbers of Table 2 would drop proportionally with numbers presented in Table 3. So you can see, if you were operating a Liebert CW106 at 80%, for instance, it would be producing airflow at a similar power level as the ARC100 running at full speed.

Other Considerations

The ASHRAE 90.1 standard on facility efficiency is being updated this year with some significant mandates for data center design. For the first time, economizers will be required for data centers of any significant size unless they fit into some fairly narrow limits. The standard still has to be adopted by local governance. But considering the ASHRAE changes have been driven by federal energy goals, it

[5] Liebert XD products do not have full variable capability. There are opportunities to shed fans with some of the XD products.

is likely there will be significant pressure to adopt quickly. Even though most IT hardware has temperature and humidity specifications that are commensurate with a wide range of economizer operation, many facilities default to ASHRAE recommended guidelines for operating conditions. The ASHRAE recommended window is quite narrow and renders a positive ROI (Return on Investment) impossible for economizers in many areas of the U.S. With an economizer, you are either going to bring fresh air into the data center and be subject to low and high humidity, or you are going to use a water economizer to cool your chilled water. Both of these strategies are challenged by ASHRAE's narrow window. There are two schools of thought,

- ASHRAE widens the window, or a significant number of data centers ignore the recommendations. In this case, more facilities may choose air economizers because of lower cap costs and wider operating range.
- Facilities will be driven to water economizers mainly due to excessive humidification associated with air economizers operating in ASHRAE's narrow range.

With the first option, fresh air is brought into the facility. Room-level distribution for the air is necessary. This economizer strategy does not work well in conjunction with row or rack coils since they are local distributors of cooling. In fact, many economizer designs route outside air through CRAC/CRAH units to use them as distribution agents even when they are not cooling the air. Row and rack coil strategies are better matched to water economizers. This should be a long term consideration in your coil selection process. If your strategy is moving toward local row or rack cooling, you may be limiting future choices in economization. While the ASHRAE mandate applies mainly to new facilities, it is written to include significant retrofits. The size/type of retrofit will be left up to your local authority.

Summary

There is much to consider when choosing a cooling coil strategy. Whatever the choice, the efficiency differences will depend more upon operating conditions than inherent differences of the coil products themselves. Your best bet is probably the product that simplifies IT deployments to your closest satisfaction. On one hand, you can choose a more expensive option, the door coil, and you get all the benefits—built in containment, no need for airflow matching, and above dew point operation—or you can choose the other end of the spectrum, a large CRAC/CRAH unit where you have to add containment and choose to operate in a more efficient manner. Just make sure you are taking everything into consideration when doing a ROI analysis and make sure you are not comparing apples to oranges. The industry is moving toward modular deployments for simplicity and repeatability. The row and rack products obviously fit this strategy, but don't count CRAC/CRAH units out of that strategy. With containment, they can achieve similar efficiency and still be fit into a modular deployment strategy. Don't forget to calculate all containment costs including the raised floor, ceiling, and containment into the ROI calculation. Economizers are likely in your future. Figure that likelihood into your long range plans; it may affect your choice in cooling coils.