

WHITE PAPER

How IT Decisions Impact Data Center Facilities: The Importance of Mutual Understanding

By
Ian Seaton



+1.888.982.7800

upsite.com

Table of Contents

Executive Summary	1
Background	2
Section 1: Specifying High ΔT Servers vs. Low ΔT Servers	2
Section 2: Specifying A1, A2, A3, or A4 Servers	5
Section 3: Specify Equipment That Breathes from Front-to-Back	8
Section 4: Specify Solid State Storage or Tape Storage	11
Section 5: Specify Cooling Unit Set Point	11
Section 6: Specify Cages That Are Compatible with Containment	14
Section 7: Airflow Management Best Practices	14
Raised-Floor Grommets and Blanking Panels	14
Cable Management	15
Server Virtualization	16
Humidity	17
Return Air Isolation	17
Implementation	17
Summary	17
About the Author	18

Table of Tables

Table 1: Comparison of Partial Free Cooling Hours at Different Server ΔT 's	4
Table 2: 2012 ASHRAE TC9.9 Data Processing Equipment Environmental Specifications	6
Table 3: Free Cooling Hours at Recommended Maximum Server Temperature and at Allowable Temperatures for Class A2, A3, and A4 Servers	7
Table 4: Chiller Plant Operating Costs at Different IT Loads and Different Set Points	12
Table 5: Indirect Evaporative Free Cooling Hours at Different Set Points	13

Table of Figures

Figure 1: CFD Model of Data Center with Uncontrolled Switch Airflow	9
Figure 2: CFD Model of Data Center with Controlled Switch Airflow	10
Figure 3: Centrifugal Chiller Operating Costs at Different Leaving Water Temperatures	12
Figure 4: CFD Model of Data Center without Floor Grommets and with Floor Grommets	15
Figure 5: Underfloor Air Passages Blocked by Excess Cable	15
Figure 6: Cable Management for a Large Switch	16
Figure 7: Server Virtualization Adoption over Time	16

How IT Decisions Impact Data Center Facilities: The Importance of Mutual Understanding

Executive Summary

Decisions and actions typically under the jurisdiction of the IT side of data center management can have a profound impact on the mechanical plant functions and resultant operating costs of the data center. By understanding these relationships, IT and facilities management should be able to form a more cooperative approach to managing the data center, resulting in a more effective and efficient operation, thereby better fulfilling often contradictory objectives.

Key considerations and findings of this paper include:

Specifying High ΔT Servers vs. Low ΔT Servers:

Higher temperature rise (ΔT) means less airflow volume per kW for cooling. This results in a lower total conditioned airflow rate and less fan energy, reducing operating cost. Historically, blade systems have had higher ΔT 's than 1U-3U rack-mount servers, though some newer Energy Star servers are closing that gap.

Specifying A1, A2, A3, or A4 Servers: Different server classes operate at different maximum temperatures, resulting in more free cooling hours, significantly reducing operating cost.

Specifying Equipment That Breathes from Front-to-Back: Not doing so violates hot and cold aisle segregation, resulting in the need for lower set points and higher air volume delivery, and therefore reduced total and redundant cooling capacity and higher operating costs.

Specifying Solid State Storage or Tape Storage:

Rate of temperature change maximums for tape storage may limit access to free cooling hours in some climate zones.

Specify Cooling Unit Set Point: Best practice is to specify maximum allowable IT equipment inlet temperature and let mechanical plant find its own level.

Specify Cages that are Compatible with

Containment: If not properly done, there will be a need for extra volume airflow and lower cooling unit set points, resulting in higher operating costs.

Employ Airflow Management Best Practices:

Adhering to airflow management best practices can reduce required airflow volume and allow you to increase cooling unit set points, lowering energy fan costs, lowering chiller fan costs, and increasing free cooling hours.

The end goal in discussing these considerations is to have an agreement between IT and facilities as to how these decisions impact the other. To ensure efficiency and optimize the data center, both teams need to work together and understand the impact that the aforementioned considerations will ultimately have on the data center environment.

Background

Specifying server, storage and telecommunications hardware, qualifying applications, maintaining virtual uptime, managing hardware and service providers, and forecasting transaction and storage capacity requirements beyond business strategic horizons ought to be job enough for IT management; however, it turns out that most of the traditional and generally-understood tasks bear heavily on the efficiency and design capacity of the mechanical plant, as well as holding the key to either optimizing or sub-optimizing the efficacy of data center architectural plans. For these reasons, we recommend that many of these decisions and routine activities be conducted in concert with facilities management and/or architectural engineering resources for new spaces to decrease the likelihood of finding yourself a year or two later painted into a corner that prevents growth, responses to competitive situations, or delivering a needed point to the organization's bottom line.

In this paper, we will explore the ramifications of some of the more common IT management responsibilities on data center mechanical plant efficiency and architectural effectiveness. This will provide a basis for IT management involving facilities management in situations with more overlap than has typically been recognized. In essence, communication and positive working relationships of both IT professionals and data center facilities managers are paramount for the mutually-beneficial outcome of an optimized data center. Additionally, certain sections of this paper address areas pertaining to specific types of data centers (e.g. Section 6 on specifying cages for containment is more germane to multi-tenant data centers), but the majority of the paper applies to all forms of data centers.

Section 1: Specifying High ΔT Servers vs. Low Symbol T Servers

- Higher ΔT typically associated with blades results in less airflow volume per kW for cooling. This results in a lower total conditioned airflow rate and less fan energy, reducing operating cost.

Specifying computer equipment for a data center might on the surface appear to be one of the most basic of IT activities, but it can also have a profound impact on the data center mechanical plant and the overall cost of operating the data center. There are many aspects to the decisions on the variety of computer servers and associated equipment that would be deployed in a data center, all of which

have some degree of impact on the efficiency of the mechanical infrastructure (specifying traditional discrete servers or blade servers is just one of those decisions). The IT decision process might typically focus on which type of server might best support virtualization or application segregation, or which might offer an easier path to technology refreshes for higher transaction speeds, or perhaps weigh the importance of I/O expansion scalability versus raw computing power; in addition, the effect on the total architecture of the data center should also be part of that decision process.

Historically, blade servers have produced higher ΔT 's than traditional rack-mount servers, popularly referred to as "pizza box" servers. That is to say, the cool supply air entering a blade chassis would exit as hotter air than would the supply air entering a pizza box server. This difference is described by the basic equation of heat transfer:

$$q = Cp \times W \times \Delta T$$

where q = amount of heat transferred

Cp = specific heat of air

ΔT = temperature rise of air across the heat source

W = mass flow

When we normalize the terms for our comfort zone of familiarity, this relationship is described as:

$$CFM = \frac{3.16 \times \text{watts}}{\Delta T}$$

where CFM = cubic feet per minute of airflow through the server

3.16 = factor for density of air at sea level in relation to °F

ΔT = temperature rise of air passing through the server in °F

Based on this relationship, a 5kW blade server chassis with sixteen (16) servers and a 35° F ΔT would draw 451.4 CFM

$$451.4 \text{ CFM} = \frac{3.16 \times 5000}{35^\circ \text{F}}$$

whereas ten 500-watt pizza box servers with a 20° F ΔT would draw 790 CFM.

$$790 \text{ CFM} = \frac{3.16 \times 5000}{20^\circ \text{F}}$$

In a data center with 1600 blades (100 chassis), the servers would consume 45,140 CFM of chilled air (100 X 451.4) as opposed to a data center with 1000 pizza box servers, which would consume 79,000 CFM of chilled air. There are several ramifications to this airflow consumption difference. First, there would be a significant cooling unit fan energy savings in the data center with higher ΔT computer equipment. For example, if we specified 30-ton computer room air handlers (CRAH) with capacity to deliver 17,000 CFM each, we would need five CRAHs to meet the 79,000 CFM requirement.¹ Naturally, those five CRAHs would be significantly over capacity for the lower airflow requirement of the higher ΔT data center, but fan affinity laws mean an even greater energy savings is achievable than might be expected in a straight linear relationship.

Those relationships are described by the equations:

$$\frac{Q_1}{Q_2} = \frac{N_1}{N_2}$$

and
$$\frac{P_1}{P_2} = \frac{N_1^3}{N_2^3}$$

- where
- Q_1 = actual airflow
 - Q_2 = maximum airflow capability
 - N_1 = actual fan speed (rpm)
 - N_2 = maximum fan speed capability
 - P_1 = numerator of power ratio
 - P_2 = denominator of power ratio

Therefore, at 100% fan speed, 100% of the rated airflow volume is delivered, and at 50% fan speed, 50% of the rated airflow volume is delivered. Likewise, at 100% fan speed, 100% of the rated fan energy is consumed, but at 50% fan speed, only 12.5% of the rated fan energy is consumed, i.e., $(1/2)^3 = 1/8$, or .125.

For the sake of discussion, then, let's assume the 30-ton CRAHs in the example have approximately 12 horsepower motors rated at 9kW. The five CRAHs can produce 85,000 CFM (5 X 17,000), but only 79,000 CFM is actually required, so the variable air volume fans (either variable frequency drive [VFD] or electronically commutated [EC or "plug"] fans) could be turned down to 93%, or 15,800 CFM each, to meet the server requirements. At this flow or rpm level, each CRAH would be drawing 7.2kW:

$$93\%^3 = 80\%$$

$$80\% \text{ of } 9\text{kW} = 7.2\text{kW}$$

With five CRAHs running 8760 hours a year, therefore, the blade data center would require 59,000 kW hours for cooling and the pizza box server data center would require 315,360 kW hours for cooling, or a 81% cooling energy savings for the data center with the servers with a higher ΔT.

While such an "apples-to-apples" comparison is useful for illustrating relative differences, such a comparison is more hypothetical than realistic. In reality, what would be more likely for the example scenarios would be that six 30-ton CRAH units might be deployed for the pizza box data center to provide an N+1 availability hedge and four of the same CRAHs would provide the same N+1 redundancy for the data center with the higher T equipment. In this case, a total cost of ownership calculation would include the difference of purchasing an additional two CRAHs and then the energy required to power those CRAHs in each scenario per year, per technology refresh cycle or per the estimated life of the data center, whichever metric is most relevant to the particular business model.

¹Note: This capacity plan is purely hypothetical just to establish a relative difference between the two server equipment scenarios. In actual practice, as will be discussed later in this paper, it could require anywhere from five CRAHs to as many as twelve or more to satisfactorily meet the 79,000 CFM requirement of the servers, depending on the effectiveness of the total data center design for minimizing stranded capacity associated with bypass and re-circulation.

By the same methodology, the blade server data center, with the same IT kW load, requires only 45,140 CFM, so each of the CRAHs could be turned down to 53% rpm, resulting in an energy requirement of only 1.35kW each ($.53^3 = .15$ and 15% of 9kW = 1.35)

In addition to directly impacting the energy required to produce and move cooling air in the data center, the decision to specify blade servers or pizza box servers can also have an effect on the performance of an economizer or free cooling facility. Free cooling is discussed in more detail later in this paper, but for the purposes of this discussion let's consider an airside economizer and a set point for 68° F supply temperature delivered into the data center. Everybody understands that as long as the outside ambient temperature is below 68° F, we can use that air to cool our data center and save the energy required to remove the heat from the data center air and return the chilled air to the data center.² Next to the energy required to power the data center's IT equipment, the single largest component of the data center operating budget is mechanical cooling, whether that is a centralized chiller plant or discrete DX CRAC

units on the data center floor. Besides eliminating this mechanical plant expense when the ambient conditions are at or below the desired supply air temperature, there is a more or less linear partial reduction of those mechanical cooling costs whenever the ambient condition is lower than the data center return temperature. In the current example, therefore, there would be some partial free cooling benefit in the pizza box server data center as long as the external ambient conditions were below 88° F and in the blade server data center as long as the external ambient conditions were below 103° F (68+20 and 68+35, respectively).

Table 1 demonstrates that this value of additional partial free cooling hours will vary greatly based on the location of the data center and, by implication, also by the operational set point.

Table 1
Comparison of Partial Free Cooling Hours at Different Server ΔT's

City	Free Cooling Hours	Partial Free Cooling Hours		Extra Hours
		At 20° F ΔT	At 35° F ΔT	At Higher ΔT
Albuquerque	5664	2704	3006	11.2%
Austin	3781	4422	4971	12.4%
Beijing	5344	3138	3416	8.9%
Boston	6684	1956	2076	6.1%
Charlotte	4950	3295	3808	15.6%
Chicago	6095	2535	2665	5.1%
Dallas	3784	4237	4970	17.3%
Denver	6653	1787	2107	17.9%
Dubai	836	4297	7608	77.1%
Frankfurt (Germany)	7478	1265	1282	1.3%
Hong Kong	1691	6524	7069	8.4%
Jacksonville	3689	4615	5071	9.9%
Las Vegas	4092	2767	4357	57.5%
Minneapolis	7874	886	886	0%
Phoenix	3015	2961	5148	73.9%
San Jose	6437	2171	2323	7.0%
Sydney	4987	3746	3773	0.7%

²Later in this paper we will discuss various set point strategies and how they relate to IT decisions and initiatives. We will use the 68° F supply air temperature in the present discussion to avoid predisposing any inherently conservative IT sensibility that we are suggesting any radical departure from the reader's comfort zones.

Just a few years ago, industry experts could confidently base analyses and energy use forecasts on 20° F as a standard ΔT for rack mount pizza box servers and 35° F as a standard ΔT for blade server systems. The playing field, however, has evolved over time. Server manufacturers continue to endeavor to produce more energy efficient servers. Much of this design energy is focused on microprocessor cores and software that slows down the computer when it is not working at high activity levels. Nevertheless, some of this work is also focused on the operating temperature specifications of internal components, resulting in reduced airflow requirements and higher allowable ΔT 's. Therefore, while the general distinction between pizza box servers and blade servers is still typically valid, it is not an absolute rule, so the specifying effort to exploit this element should require the server vendor to provide either ΔT information or CFM per kW ratios. Only by considering this element of a server specification can the server acquisition and deployment activity treat the data center as a totally integrated eco-system.

Section 2: Specifying A1, A2, A3, or A4 Servers

- Different server classes operate at different maximum temperatures, resulting in more free cooling hours, significantly reducing operating cost.
- Different maximum server temperatures result in different chiller operating conditions and may mean no chiller is required at all, thus dramatically reducing operating cost and potentially significantly reducing the cost to construct new data centers.

In order to at least begin the process of building a fully integrated data center that functions as an eco-system wherein all the various elements are mutually supportive to optimize the IT, mechanical and electrical

functions, decisions about servers should include more variables than the previously-described impact of rack-mount servers versus blade systems and all the associated application compatibility and I/O versus raw power considerations. Servers can also be categorized by their operating temperature boundary requirements, and these categories can have a dramatic effect on the operating energy budget of a data center, the total cost of ownership of a data center operation, and the mechanical robustness of the data center. These categories were formalized by the ASHRAE Technical Committee 9.9, IT sub-committee in a 2011 white paper, and then in the third edition of its handbook on environmental guidelines for data center computer equipment (summarized in Table 2).

However, before exploring the ramifications of IT decisions from the data in this table, we should first be clear on the source of this information as a pre-emptive strike against concerns that the server manufacturers would not cooperate with any of the recommendations and suggestions. The developers of these guidelines, i.e., the members of the ASHRAE TC9.9 Mission Critical Facilities IT sub-committee are, in fact, the server OEMs – research fellows, emeritus scientists, and graduate-degreed engineers representing those companies comprising over 70% of the global server market³ as well as the bulk of the server microprocessor industry. Nothing is being proposed beyond the published and well-understood capabilities of the data processing equipment that typically finds its way into data centers and computer rooms. With that pre-emptive clarification, then, we can explore these additional procurement options facing today's IT managers and what they mean for the overall effectiveness and efficiency of the data center.

³ IDC Worldwide Quarterly Server Tracker, December 2014

Table 2
2012 ASHRAE TC9.9 Data Processing Equipment Environmental Specifications⁴

Classes	Equipment Environmental Specifications							
	Product Operations					Product Power Off		
	Dry-Bulb Temperature (°C)	Humidity Range, non-condensing	Maximum Dew Point (°C)	Maximum Elevation	Maximum Rate of Change (°C/hr)	Dry-Bulb Temperature (°C)	Relative Humidity (%)	Maximum Dew Point (°C)
Recommended								
<i>(Applies to all A classes; individual data centers can choose to expand this range based upon the analysis described in this document)</i>								
A1 to A4	18 to 27	5.5°C DP to 60% RH and 15°C DP						
Allowable								
A1	15 to 32	20% to 80% RH	17	3050	5/20	5 to 45	8 to 80	27
A2	10 to 35	20% to 80% RH	21	3050	5/20	5 to 45	8 to 80	27
A3	5 to 40	-12°C DP and 8% to 85% RH	24	3050	5/20	5 to 45	8 to 85	27
A4	5 to 45	-12°C DP and 8% to 90% RH	24	3050	5/20	5 to 45	8 to 90	27
B	5 to 35	8% to 80% RH	28	3050	N/A	5 to 45	8 to 80	29
C	5 to 40	8% to 80% RH	28	3050	N/A	5 to 45	8 to 80	29

The 2011 ASHRAE environmental guidelines maintained the recommended server inlet temperature range established in 2008 of 64.4°F to 80.6°F, but expanded the allowable temperature range and corrected a problem it had introduced in applying the allowable range. The ASHRAE paper that introduced the first expanded temperature range advised that it was OK to operate within the allowable range for short periods of time, without ever defining a “short period of time.”⁵ ASHRAE’s 2011 white paper and third edition of the handbook provide clear direction on how to systematically determine an acceptable short period of time for any data center. What follows is an extremely abridged synopsis of the steps for determining that “short period;” the

authors encourage the reader to consult the handbook and include this analysis as part of the normal pre-specification/acquisition due diligence.

The ASHRAE IT committee members accomplished an unprecedented level of cooperation by establishing what they called an “X” factor, which represented an expected reliability level that would be achieved by operating a data center with every server ingesting 68°F air 24/7 over the life of the data center or over the life of a technology refresh. They then agreed on percentage levels up or down from that baseline based on how many hours servers would receive air either below or above that “X” factor baseline. The handbook then presents case study examples on how

⁴Thermal Guidelines for Data Processing Equipment, 3rd Edition, ASHRAE, 2012

⁵Environmental Guidelines for Datacom Equipment: Expanding the Recommended Environmental Envelope” ASHRAE, August, 2008, p.2.

to use the guidelines. In one example, it describes a data center in Chicago using air-side economization free cooling that allowed the data center temperature to fluctuate, following Mother Nature’s lead, within a particular allowable range. Temperatures below the low threshold were accommodated merely by re-circulating return air to keep the supply warm enough, and temperatures above the high threshold were addressed with standard mechanical cooling. In the Chicago example, the server equipment reliability actually improved by 1%.

Interpretation of the results of this exercise merits a word of caution. For example, if the calculations using hourly temperature bin data for the data center location indicated that allowing the data center temperature to fluctuate based on anticipated weather conditions would produce a 5% increase in server failures, what does that actually mean? If the data center had 1000 servers and the experience from the particular vendor translated into an expected 0.5% failure rate, that would mean an expectation

for five failed servers somewhere during their lives. A 5% increase over five servers means that allowing the data center temperatures to fluctuate up and down within one of the allowable ranges would result in increased failures of ¼ of a server, or 1 out of 1000 over four years, which may actually be longer than the technology refresh for some data centers. Determining the scale of this perspective and what might be acceptable should be a preliminary step before making decisions about server classes that can have profound effects on both the design and efficiency of the data center mechanical plant.

For the Celsius-challenged, the allowable range for Class 1 servers from Table 3 is 59° – 90° F, for Class 2 servers is 50° – 95° F, for Class 3 servers is 41° – 104° F and for Class 4 servers is 41° – 113° F. It will still be difficult to find manufacturer designations for these classes on the standard marketing or user documentation; however, these ranges will correspond to the operating temperature ranges that are readily available on all manufacturers’ standard documentation.

Table 3

Free Cooling Hours at Recommended Maximum Server Temperature and at Allowable Temperatures for Class A2, A3, and A4 Servers

City	Airside Economizer				Indirect Evaporative			
	Rec	Class A2	Class A3	Class A4	Rec	Class A2	Class A3	Class A4
Albuquerque	7639	8701	8760	8760	8760	8760	8760	8760
Austin	7120	8733	8760	8760	6481	8758	8760	8760
Beijing	7489	8717	8760	8760	7761	8760	8760	8760
Boston	8291	8747	8760	8760	8600	8760	8760	8760
Charlotte	8028	8760	8760	8760	8479	8760	8760	8760
Chicago	7339	8676	8760	8760	8268	8756	8760	8760
Dallas	6697	8617	8759	8760	6972	8758	8760	8760
Denver	7941	8699	8760	8760	8760	8760	8760	8760
Dubai	3323	6708	8445	8755	5414	8758	8760	8760
Frankfurt (Germany)	8580	8758	8760	8760	8758	8760	8760	8760
Hong Kong	4977	8744	8760	8760	5014	8760	8760	8760
Jacksonville	7104	8720	8760	8760	6486	8760	8760	8760
Las Vegas	5749	7698	8522	8752	8674	8681	8760	8760
Minneapolis	8020	8744	8760	8760	8586	8760	8760	8760
Phoenix	4635	7203	8244	8747	8078	8760	8760	8760
San Jose	8246	8724	8760	8760	8745	8760	8760	8760
Sydney	8609	8758	8760	8760	8728	8755	8760	8760

Just for perspective, the very significant majority of servers sold at the time of writing this paper fall into the Class 2 range. Class 4 servers today are typically the extremely ruggedized servers for military or other harsh field applications. Class 3 servers fall somewhere between those extreme ruggedized servers, which typically come at a significant price premium, and what we find in most commercial critical facilities. As data center operators become more familiar and comfortable with incorporating this type of analysis as part of the standard specification and acquisition process, we can expect more attractive pricing on servers with these temperature specifications but without the rest of the ruggedized architecture. All that notwithstanding, those price premiums should not, as a matter of course, disqualify these servers from consideration, particularly considering the impact on total cost of ownership available in many situations.

Table 3 shows the potential opportunity for free cooling at allowable temperature maximum thresholds for different classes of servers. Airside economizer hours are based on dry bulb temperatures for the indicated cities and indirect evaporative cooling hours are based on wet bulb temperatures for those cities. The other primary difference between the two example scenarios is that airside economization allows outside air into the data center, whereas indirect evaporative cooling, by definition, conducts heat transfer through a heat exchanger that keeps outside air separated from inside data center air. Between the two types of free cooling, the table clearly shows that data centers in Albuquerque, Boston, Chicago, Beijing, Denver, Frankfurt, Hong Kong, Jacksonville, Minneapolis, Phoenix, and San Jose could cool data centers with Class A2 servers (typical commercial servers today) all year without ever using any refrigerant mechanical cooling. In addition, Austin, Charlotte, Dubai, Las Vegas and Sydney could cool Class A3 servers all year with one form or the other of free cooling. While this is obviously only the first step of the analysis, the fewer hours between the ASHRAE-recommended maximum temperature and the allowable temperature, the more likely the "X" factor reliability forecast will justify allowing the temperature fluctuations. So what does that mean?

One path suggested by such an analysis is a data center design completely devoid of any chiller or other refrigerant-based mechanical cooling system. Such a design not only saves on the cost of mechanical cooling, which is the single largest operational cost of a data center after energy required to power the IT equipment, but it could also save on the capital

investment of the mechanical cooling plant and the real estate space required for its deployment.

On the other hand, a budget analysis might indicate that Class A3 servers over the three to six technology refreshes that might be planned for a data center life might actually be a higher investment than building a chiller plant and associated systems; in this case, the return on investment would be the operational savings for using free cooling for 80% or 90% of the year, or whatever the bin data indicate should be expected. Even in this scenario, it's possible that a capital savings could be achieved by using the free cooling facility and the allowable hours bucket as the redundant back-up to the chiller, and thereby provide relief from the capital investment for redundant chillers and air handlers or CRAC units. Even in cases where the traditional mechanical plant is designed with an N+1 level of redundancy, the free cooling capability and the allowable temperature hours bucket can provide another level of redundancy, raising the cooling system from a basic level of availability protection to a more clearly fully fault tolerant and simultaneously maintainable capability, for a level of robustness that otherwise might not be affordable.

By this time, after looking only at the potential differences between standard rack mount servers and blade systems and the opportunities associated with the new ASHRAE TC9.9 server classes, it should be clear that server acquisition choices have a bigger impact on the mechanical operating efficiency of a data center than do choices of CRAC units, heat rejection strategies, chillers, and those elements typically under the authority of the functional areas responsible for the data center energy budget. And there are many more decisions belonging to the IT functional area that directly impact both the overall operational efficiency of the data center as well as the effectiveness and robustness of its design.

Section 3: Specify Equipment that Breathes from Front-to-Back

- Not doing so violates hot and cold aisle segregation, resulting in the need for lower set points and higher air volume delivery, and therefore reduced total and redundant cooling capacity and higher operating costs.

Servers are not the only equipment IT management is tasked with specifying for the data center: the number of network switches may not approach the number of servers, but they are still critical to the effective functioning of the IT enterprise. The selection


process for switches will typically consider how many ports are required, whether it will be a core switch or distribution switch with layer 3 or router functionality, whether it can be unmanaged or must be managed, whether speed and latency or raw volume is most important, and what might be required for power over Ethernet and redundancy. In addition, switches are typically sourced with some attention to scalability, plans for virtualization, ease and speed of application deployment and speed of application access. Investigating and sourcing switches should be closely tied to the overall business needs.

To the degree that efficiency and profitability also fall into the area of core business goals, the effect of the switches on the overall effective airflow management of a data center should also be part of that overall criteria checklist. As far as the mechanical health of the data center is concerned, all equipment should be rack-mountable and breathe from front to rear. This criterion is a recognized best practice from sources as diverse as the European Code of Conduct for Data Centers, BICSI-002, and ASHRAE TC9.9. If only it were so simple.

Nevertheless, even though there may be a small fraction of switches in a data center compared to the


number of servers and storage devices, the effect of non-standard airflow can be staggering. An example of a 3500 square foot 514kW computer room is illustrative of the effect of uncontrolled non-standard-breathing network hardware. Figure 1 below shows the effect of this uncontrolled airflow. There are six racks in the point of presence entrance room with only 3kW per rack, but the equipment in one rack is ingesting inlet temperatures ranging from 73° F to 90° F and one 6kW core switch rack is ingesting air ranging from 56° F to 85° F. While those temperatures were actually allowed by the manufacturers' specifications, they exceed the owner's internal SLA's, but more importantly they undermined the otherwise efficient design of this space. All the server cabinets were equipped with exhaust chimneys coupled to a suspended ceiling return space and the supply air was ducted directly into all the cold aisles. In addition, ceiling grates were located around the uncontained cabinets and racks to capture the heated return air. Nevertheless, in order to minimize the number of hot spots and the severity of those hot spots, the cooling system was configured with a 72° F set point, resulting in 54° F air typically delivered through the overhead ducts. In addition, the cooling system needed to deliver 82,500 CFM to meet the 67,900 CFM demand of the IT equipment.

Figure 1
CFD Model of Data Center with Uncontrolled Switch Airflow



Contrast those results with a model of the same data center with all the return air effectively contained. Side-breathing switches were re-deployed in cabinets that redirected the airflow into a front-to-rear topology, and those cabinets were then captured in a hot aisle containment structure. In the entrance room, the side-to-front and rear-to-front switches were removed from open two-post racks and re-deployed in cabinets equipped with rack-mounted boxes, which redirected the airflow into a standard front-to-rear topology, and then those cabinets were equipped with vertical exhaust ducts.

All relevant best practices and standards dictate that proper airflow management can only be achieved with equipment that breathes from front to rear. Unfortunately, other application criteria often result in selections of non-conforming equipment. Fortunately, there are ways to change any equipment into exhibiting an output behavior equivalent to front-to-rear airflow. For large switches that breathe side-to-side, there are now wider cabinets available with internal features that will redirect that airflow into conformance with the rest of the data center. For smaller-profile switches that may have front-to-side airflow, or side-to-front or side-to-rear or even



In addition, because these changes eliminated all the hot spots throughout the room, the need for the overhead duct delivery system was eliminated, removing that cost from the total capital plan. Figure 2 illustrates the results of these changes. Not only is the temperature more consistent and hot spots eliminated, the supply temperature was raised from 54° F to 75° F, resulting in energy savings of 38% for the chiller plant. Furthermore, the airflow requirement to maintain all equipment inlet temperatures between 75°-76° F dropped from 82,500 CFM to 72,000 CFM, resulting in a 33.5% CRAH fan energy savings, based on the non-linear fan law relationships discussed earlier in this paper.

rear-to-front, there are rack mount boxes available that will redirect any of these airflow patterns into conformance with the rest of the data center. These boxes will typically consume one or two extra rack mount units, but they will maintain the integrity of an otherwise well-designed and -executed space. Some storage equipment will present similar challenges to the data center manager; however, some form of correction is available and should be utilized to either maximize the efficiency of the data center mechanical plant or to increase the effective capacity of that mechanical plant.

Section 4: Specify Solid State Storage or Tape Storage

- Dramatic differences in required rates of temperature change could impact access to free cooling.
- Different RH requirements change access to free cooling and requirement for humidity management.

Decisions about storage architectures are another IT management responsibility with implications on the effectiveness and efficiency of the mechanical plant. While cost has inhibited widespread adoption of solid state storage,⁶ the impending exponential growth in storage capacity requirements may speed up the tipping point on that adoption rate. However, there are also mechanical plant operational differences between tape storage and solid state storage. The previously-cited ASHRAE TC9.9 environmental guidelines for data processing equipment state a significant difference between the allowable rates of temperature change for tape storage versus solid state storage. While supply temperature rate of change for tape storage is limited to 9°F per hour, that same parameter for solid state storage is 36°F per hour. On the surface, that distinction may not seem that critical because it may be difficult to imagine 9°F, much less 36°F temperature swings in a data center. However, the specification is for “rate of change,” not total absolute change; therefore, a 2°F change in 10 minutes equates to a rate of change of 12°F per hour, and would therefore exceed this maximum requirement.

In situations discussed earlier in this paper where a data center could be designed and constructed without any refrigerant mechanical plant or where an economizer system might perform either an N+1 or 2N availability guarantee for a data center, it is easily conceivable that the data center could potentially be exposed to temperature fluctuations exceeding this threshold. Therefore, the potential 3:1 acquisition and maintenance premium associated with solid state storage should be evaluated in terms of the cost avoidance of both a chiller acquisition and its installation, as well as its operating cost over the life of the data center.

⁶James Maguire, “Leading Data Storage Trends: Solid State and Software Defined,” *Enterprise Storage Forum*, April 30, 2014.

Section 5: Specify Cooling Unit Set Point

- Best practice is to specify maximum allowable IT equipment inlet temperature and let mechanical plant find its own level.
- Managing temperature by thermostat set point frequently results in mechanical plant wasted energy and cycling or heating by cooling equipment.

While determining a data center set point may sound like a facilities decision, it is often dictated by IT management, leading to the normal round of “meat locker” jokes at data center conferences around the country. The plethora of conference papers and industry e-newsletters notwithstanding, it is still common to see 72°F set points in a large proportion of data centers today, and lower set points in the range of 68°F are still not uncommon. These set point decisions can be driven by inherent IT conservatism, precedent (“that’s the way we’ve always done it”), or by response to real or imagined hot spot problems. This latter hot spot mitigation motivation can be problematic. Remember that most data center cooling units will drop the temperature of air meeting or exceeding the set point by approximately 18°F, thus a 68°F set point could conceivably produce 50°F supply air, which could be cool enough to get us into dew point problems. At 50°F, saturation (or 100% RH or condensation) is reached with 55 grains of moisture per pound of dry air, a condition which would be met at any of the following data center control settings:

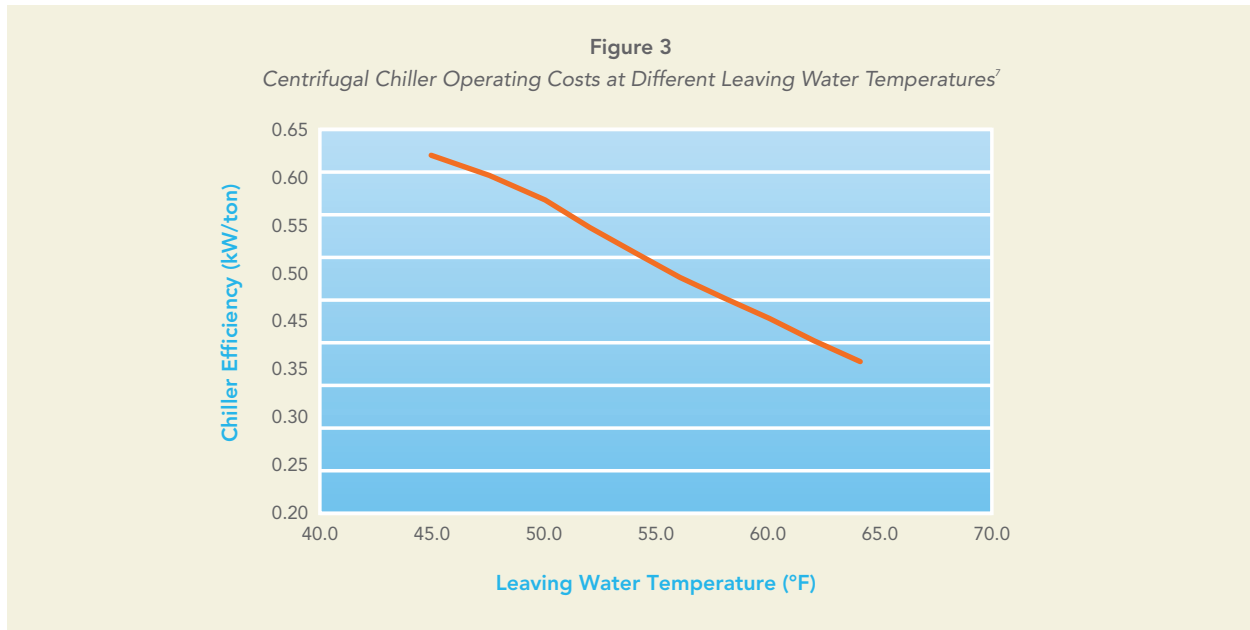
60% RH @ 65°F

50% RH @ 70°F

45% RH @ 75°F

36% RH @ 80°F

In other words, at 50°F, the cooling coils would start removing water from the air in the form of condensation. In the mechanical side of the data center, we call this latent cooling, which subtracts from the sensible cooling capacity of our cooling units. Therefore, it’s conceivable that the lower set point could actually reduce the cooling capacity of the data center mechanical plant. Related to this ironic outcome is a potential result that may stem from fighting hot spots by adding extra cooling units. While the decision of how to solve a cooling problem is not typically an IT decision, it is noteworthy because, as might be the case for reducing set points, the extra



cooling units may produce the similarly-ironic result of reduced cooling capacity and exacerbated hot spot issues. What can happen is that a large over-capacity of supply air can end up short-cycling to the cooling unit without picking up any heat, resulting in return air below the set point, which will cycle off the cooling coils and just recycle the cooler air back into the data center. Therefore, rather than 55° F or 60° F air being pushed into the data center by those fans, that air may be 69° F or 70° F and actually raise the temperature of that part of the data center experiencing the hot spots.

More importantly, arbitrarily-established low set points drive up the cost of running the data center cooling plant. Figure 3 shows kilowatt energy consumption of a highly efficient centrifugal chiller plant at different leaving water temperatures. Note that this level of absolute efficiency may not be achieved by older chillers or plants that have not been optimized, but the relative differences of the slope are instructive and many data centers can reasonably expect a steeper savings line. This line graph offers a conservative and defensible estimate of chiller energy use, and these values were used in calculating the actual energy costs in Table 4.

Table 4
Chiller Plant Operating Costs at Different IT Loads and Different Set Points⁸

IT Load	Cooling Unit Set Point				Supply Set Point
	68° F	72° F	75° F	80° F	75° F
500kW	\$105,908	\$101,835	\$96,132	\$86,356	\$61,916
1 MW	\$211,817	\$203,670	\$192,265	\$172,712	\$123,831
2MW	\$423,634	\$407,340	\$384,529	\$345,424	\$247,663
10MW	\$2,118,168	\$2,036,700	\$1,922,645	\$1,727,122	\$1,238,314

⁷ASHRAE Data Center Design and Operation Book #6: **Best Practices for Datacom Facility Energy Efficiency**, 2nd edition, 2009, p.40
⁸NB: Chiller tons based on 130% of IT load to account for battery room and electrical room cooling and capacity head room.

While Table 5 clearly shows the free cooling benefits associated with good set point management, it also shows for geographies such as Austin, TX; Dallas, TX; Dubai; Hong Kong; Jacksonville, FL; and Sydney, Australia, that with low set points there would not be access to enough free cooling hours to justify the expenditure for free cooling, whereas with good set point management, that expenditure is easily justified by one- to two-year paybacks. Therefore, IT management decisions regarding data center set points not only impact the overall operating efficiency of the data center, but the very architectural shape of the data center itself.

Table 5
Indirect Evaporative Free Cooling Hours at Different Set Points^a

City	Cooling Unit Set Point				Supply Set Point
	68° F	72° F	75° F	80° F	75° F
Albuquerque	3845	4754	5253	6143	8755
Austin	1315	1823	2210	2718	4865
Beijing	3785	4226	4558	5102	6704
Boston	3845	4418	4754	5458	8057
Chicago	3703	4230	4653	5259	7685
Charlotte	2203	2741	3160	3939	6915
Dallas	1648	2091	2491	2980	5588
Denver	4367	5149	5764	6579	8759
Dubai	0	4	28	267	3827
Frankfurt	2833	3967	4669	5963	8640
Hong Kong	45	184	337	702	3790
Jacksonville	677	1127	1501	2107	4996
Las Vegas	2380	3369	4299	5897	8090
Minneapolis	4057	4526	4890	5613	7874
Phoenix	947	1741	2535	4050	7036
San Jose	4469	4826	5351	6459	8672
Sydney	149	502	1184	2780	7616

^aNB: Free cooling hours based on hourly bin data for noted cities from 2007 and 2008. Figures will likely vary with more current data or with a ten-year sample, but relative differences between different set points are adequate for preliminary planning purposes.

Section 6: Specify Cages that Are Compatible with Containment

- Cages and containment don't have to be mutually exclusive.
- If the cage is not compatible with containment, there will be a need for extra volume airflow and lower cooling unit set points, resulting in higher operating costs

Customer/client cages in colocation data centers have historically resulted in sub-optimized airflow management, but this result isn't as necessary as most practitioners seem to believe. The basic elements of airflow management can be accomplished in cages, such as use of filler panels in unused rack spaces and floor grommets for any floor tile openings outside of cold aisles. However, the big-ticket tactic, as illustrated in the earlier discussion on network switches, is some form of air containment or separation, whether that be hot aisle containment, cold aisle containment or cabinet containment in the form of cabinet chimneys or vertical exhaust ducts; this level of airflow management has traditionally been precluded by the use of cages to separate different customer spaces.

In the earlier discussion related to Figure 2, on bringing six low-density entrance room racks and six medium-density core switch cabinets into the level of containment already achieved by the rest of the data center, we saw a 38% reduction in chiller plant energy cost and a 33.5% reduction in CRAH fan energy cost. Those results were achieved by completing the containment of 12 cabinets out of a total of 79 cabinets in this room, indicating that less than 1/6 of the cabinets, or only 11% of the total IT heat load, can have that large of an impact on the total cost of cooling a data center. If the deployment of cages has precluded the use of any containment in a particular data center, then the opportunity for total operational cost reduction can easily be double or more that shown in this example.

While cages and containment can be problematic, they do not need to be mutually exclusive. Cages with openings in the top can easily accommodate chimney cabinets. While aisle containment will require a little more creativity, using a solution such as Upsite's AisleLok® Modular Containment can help work around the barriers. The hot aisle is secured via doors with coded electronic locks. The cold aisle is secured with a sliding cage with access limited to the space renter and the site owner. The overhead power busway and network ladder rack enter the common aisle chimney through sealed penetrations. Although it

may change the appearance of a cage a bit, the basic customer-centric objectives can be achieved without compromising the integrity of the airflow management separation objective and resultant operating cost savings, whether those fall into the bottom line or become pricing enticements for new customers.

Section 7: Airflow Management Best Practices

- Adhering to airflow management best practices can reduce required airflow volume and allow you to increase cooling unit set points, lowering energy fan costs, lowering chiller energy costs, and increasing free cooling hours


Raised-Floor Grommets and Blanking Panels

At this juncture in the ongoing evolution of data center management, there should be no need to point out the value of plugging airflow holes with rack-mount blanking panels and floor tile cut-out grommets. They are a part of this discussion because they remain two of those products, services or design options that deliver efficiency benefits to the facilities side of the house but are part of the IT management procurement model. The very earliest discussions and publications on data center thermal management, back at the advent of hot aisle/cold aisle organization, included recommendations for plugging all the holes that provided paths for bypass airflow (short-cycling IT heat load) and re-circulation airflow (equipment ingesting heated return air). Originally the value of filler panels and floor grommets was touted as a hot spot solution as in the example in Figure 4, in which the addition of floor grommets in the cable cut-outs in the floor in the rear of all the cabinets resulted in a 90% increase in average airflow through all the perforated floor tiles and grates properly located in cold aisles.

However, very early on, this recapturing of stranded cooling capacity became an efficiency producer as the industry learned that the same IT load could be cooled with significantly less air volume and warmer temperatures when the effects of re-circulation and bypass airflow were greatly mitigated. In addition to floor grommets and filler panels, another element of this hole-plugging imperative is to also ensure the sealing of the area around the perimeter of the rack mounting space in a cabinet, i.e., the space between the front mounting rails and the cabinet frame or side panels. Some cabinets will close this area off by either integrating the mounting rails into the frame or by allowing an option for the rails to be located directly against the frame as a user-selected installation option. A word of caution about such configurations

Figure 4

CFD Model of Data Center without Floor Grommets and with Floor Grommets



Data Center without Floor Grommets

Data Center with Floor Grommets

is warranted: handles or other server front panel obstructions may not fit inside of the cabinet front door if the rails are fully forward. Regardless, accounting for this bypass and re-circulation path needs to be part of the IT procurement due diligence for server cabinets.

Cable Management

Cable management falls squarely under the responsibility of IT management of the data center, even if in larger organizations there is a network administration group that is still within the IT organization. Keeping network cables and power cables neatly dressed and labeled has become more and more critical for the effective management of moves, adds and changes (MACs) as density and application complexity have proliferated. However, this element also bears on the effective thermal management of the data center. Most obviously, when the underfloor space is being shared between airflow delivery and cable distribution, it is critical to minimize any impediments to airflow. The general rules of thumb for underfloor cable management are to align basket trays parallel to airflow direction and to maintain a strict discipline in removing all unused cables. At one time or another, we have all seen a horror story such as that illustrated in Figure 5.

There is more to the cable management impact on airflow than the underfloor space. Cable inside the cabinet can impede exhausting hot air out of the rear




Figure 5

Underfloor Air Passages Blocked by Excess Cable

of equipment cabinets, resulting in re-circulated hot air that increases the equipment inlet air temperature, possibly creating hot spots or widening the variation of inlet temperatures throughout the data center. These issues in turn can thereby drive down set points and thus create energy expense penalties from less efficient chillers, reduced access to free cooling hours, and likely increased fan energy. These types of obstructions can also affect the performance of server fans.¹⁰

¹⁰Moss, David, "IT Equipment Response to External Pressure," A Dell Technical White Paper, 2009

Large network switches contain very high port counts with resulting massive cable bundles, as shown in Figure 6. Not only do these cables need to be dressed out well, they will often create the need for either a wider cabinet to allow space for airflow into air inlets on the side of the switch, or some means of directing air into those side inlets around the cable bundle, or both. Finally, moving cable distribution from under the floor to overhead can help with underfloor air distribution; however, it is always advisable to locate cable pathways closer to the equipment racks than to the ceiling, to allow return air paths to congregate closer to the ceiling and discourage extra re-circulation in the data center.





Figure 6
Cable Management for a Large Switch

Server Virtualization

Server virtualization has long since passed the tipping point from buzzword to mainstream behavior. A recent study by Austin’s Spiceworks of the SMB market (summarized in Figure 7) clearly indicates the scale of this adoption rate. Virtualization alone has an obvious impact on the efficiency of the data center mechanical plant; however, those benefits are only partial until unused servers are actually taken out of service and the resultant obsolete cabling is removed from the data center, where it can provide an impediment to effective airflow.



¹¹Spiceworks, “State of SMB IT 1H 2013: Semi-Annual Report on Small and Midsize Business Technology Plans and Purchase Intent,” May 29, 2013

Humidity

While humidity specifications are executed by facilities management, they are typically dictated by IT management. If these specifications are derived from manufacturers' user documentation or guidelines from reliable sources such as ASHRAE TC9.9, then, if properly executed, humidity in the data center will be managed both effectively and efficiently. However, if those specifications arise from historic tribal knowledge, they will likely negatively impact the overall bottom line of the data center operation. Today's recommended environmental guidelines are written around dew point with a maximum 60% relative humidity, as opposed to previous specifications which typically established a range such as 40%–55% RH.

The key strategy with humidity is to first take advantage of all the airflow management strategies and tactics previously discussed in this paper to allow higher supply and return temperatures. The high temperatures will carry higher grain counts of moisture well above dew point to eliminate latent cooling penalty and maximize the sensible cooling capacity of the mechanical plant. At these higher temperatures, the allowable humidity range of 20%-80% RH become practical because the servers themselves will facilitate keeping the data center temperature above dew point. As a final word of caution, using CRAHs or CRACs as heaters to remove humidity from the data center and using steam generators to add humidity to the data center are excluded as acceptable practice by all data center standards, and are beginning to be outlawed by state and municipal building and energy codes. For good reason, we are at the tipping point for outgrowing those wasteful practices.

Return Air Isolation

Finally, the central airflow management best practice is to maximize the separation of the return air from the IT equipment air intakes. Aspects of this have been previously discussed in terms of hot aisle/cold aisle arrangement, chimney cabinets and containment aisles. Most of the hardware required for these best practices comes from IT procurement decisions, though implementation may rely on

facilities cooperation. Other means for promoting this separation include open ceiling grates in hot aisles and using the space above a suspended ceiling for a return air path, or aligning a smaller computer room with one row of server racks on a raised floor facing away from the cooling unit(s) for a very short, very direct return air path. An alternative to this small room approach is the use of a slab floor with up-flow cooling units ducted over the top of the row of server racks with the same short, direct return air path.

Implementation

In planning the design for a new data center space, all these elements of best practices should be at the foundation of any plan. However, if the project is focused on improving the performance of an existing space, an orderly approach to implementation will usually produce better results than trying to do everything at once or randomly attacking obvious problem areas.

Following the path Upsite Technologies has called the 4 R's¹² can reduce the likelihood of re-do's and the need to fix problems that have become bigger than they might have otherwise, had contributing factors been fixed first. Generally, the 4 R path involves first sealing the raised floor and then addressing the racks, then the rows, and then the room. There are steps and considerations for each of the 4 R's.

Summary

In conclusion, it should now be clear that the role of IT management in establishing mechanical efficiency opportunities in the data center and contributing opportunities for more efficient architectural design options is critical to the success of both the IT mission and the facility mission. Whether it is specifying servers, or qualifying and recommending applications, or establishing service-level agreement environmental thresholds, all those activities are critical ingredients in defining the overall operating efficiency potential of the data center. By being aware of these opportunities, IT management can work more closely with architectural and mechanical functional areas to create a data center eco-system more beneficial to the organization's overall business agenda.

¹²Montoya, Isaac, "4 Steps to Optimizing Computer Room Airflow Management," June 4, 2014
<http://www.upsite.com/blog/4-steps-optimizing-computer-room-airflow-management/>

About the Author

Ian Seaton retired in 2013 with 33 years of industry experience. He was an editor of BICSI-002 and remains a working group leader and editor. He is a patent holder for airflow performance algorithms and provides that expertise on The Green Grid's Utility



Task Force Airflow Energy Calculator Committee. He spent 16 years at Chatsworth Products, Inc., initiating its cabinet and thermal management solutions. Mr. Seaton is a past corresponding member of ASHRAE TC9.9 and has authored and taught numerous BICSI- and AIA-accredited courses. He has presented at over 100 technical conferences in eight countries and has published a dozen technical papers. He currently serves as a member of the Technical Advisory Group of Upsite Technologies.