The Unexpected Impact of Raising Data Center Temperatures

White Paper 221

Revision 0

by Wendy Torell
Kevin Brown
Victor Avelar

Executive summary

Raising IT inlet temperatures is a common recommendation given to data center operators as a strategy to improve data center efficiency. While it is true that raising the temperature does result in more economizer hours, it does not always have a positive impact on the data center overall.

In this paper, we provide a cost (capex & energy) analysis of a data center to demonstrate the importance of evaluating the data center holistically, inclusive of the IT equipment energy. The impact of raising temperatures on server failures is also discussed.

APC by Schneider Electric White Papers are now part of the Schneider Electric white paper library produced by Schneider Electric's **Data Center Science Center**
DCSC@Schneider-Electric.com

Introduction

There has been an ongoing industry effort to raise IT operating temperatures initiated by a revision to ASHRAE standard TC9.9 released in 2011. ASHRAE's recommendation is intended to lower data center energy consumption by expanding the number of hours of 'free cooling'.

Despite this industry effort, many data centers still operate at temperatures at or below 21°C (70°F). While there is an arguably conservative culture in the industry (avoid downtime, stick with what works, etc.) we were curious as to why more data centers were not raising IT inlet temperatures. If there were great savings to be had, one would think that more data centers would take advantage of these savings even given the conservative nature of the industry.

Through our research, we determined that there are generally two questions data center managers ask that prevent them from raising their data center inlet temperature:

- How much energy can I save by increasing my IT temperature?
- Will raising temperatures impact the reliability of my IT gear?

To answer these questions, it is important to look at the data center holistically as the system dynamics are complex – the energy consumption of some systems decrease while others increase. Consider a packaged chiller design: When the IT temperature set point is increased, the chiller energy decreases for two reasons; the data center can operate in economizer mode(s) for a larger portion of the year, and the chiller efficiency increases. But this isn't the entire picture. Although the chiller energy decreases, the following also occurs:

- The dry cooler (which operates in economizer mode instead of the chiller) energy increases because the number of economizer hours increases.
- Server energy increases because airflow (CFM) requirements increase as temperature rises. (see **sidebar**)
- Computer room air handler (CRAH) fans speed up to support the higher server CFM requirement, which means greater CRAH energy consumption.
- If not already oversized to accommodate the additional airflow, more CRAHs are needed to match the higher server fan CFM requirements. This means additional capital expense.

Why do server fans ramp up?

The purpose of server fans is to cool the components inside the server chassis. The most important of these components is the CPU chips which can reach temperatures upwards of 90°C (194°F). As the IT inlet air temperature increases, so will the CPU temperature. This typically triggers server fans to increase airflow in an effort to reduce the CPU temperature. This increase in airflow consequentially increases server energy consumption.

Figure 1 illustrates these countering effects. In this paper, we walk through an analysis of a data center with a packaged chiller cooling architecture, to demonstrate how location and server fan behavior have a significant influence on the potential savings (or cost penalties) when we increase IT inlet air temperature set points.

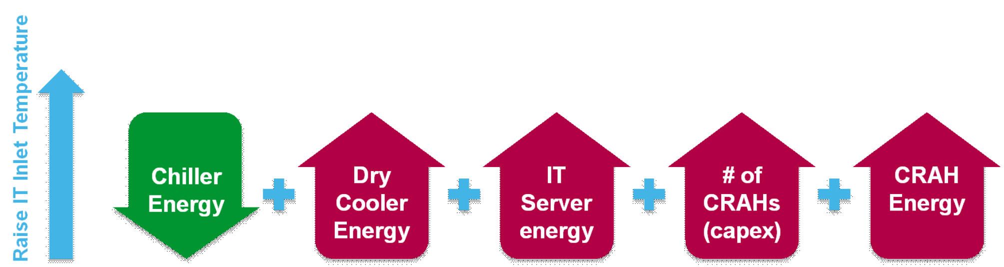

Figure 1

System dynamics are complex. Need to evaluate the data center holistically

We also look at the implications of fixing the temperature (at a higher point) vs. allowing the data center temperatures to float within a defined range, as the outdoor temperature fluctuates. Lastly, we look at a scenario where an existing data center is oversized (50%), to illustrate the impact that percent load has on these results.

Analysis of data center

The impact of raising IT temperatures on energy consumption can vary significantly depending on the cooling architecture, the climate, IT fan speed, and percent IT load. In this paper, we chose one architecture[1] and modeled it in varying climates, as an example to demonstrate the complex nature of data centers, and to illustrate the importance of understanding the risks vs. rewards before operational changes are made.

Architecture analyzed

For this analysis, we chose what we believe to be a very common cooling architecture deployed in data centers today – a packaged air-cooled chiller with economizer (**Figure 2**). The dry cooler, utilized during economizer mode, is a heat exchanger that directly cools the data center chilled water when the outside air conditions are within specified set points. Pumps move the chilled water through the dry cooler where the cold outside air cools the chilled water that supplies the CRAHs.

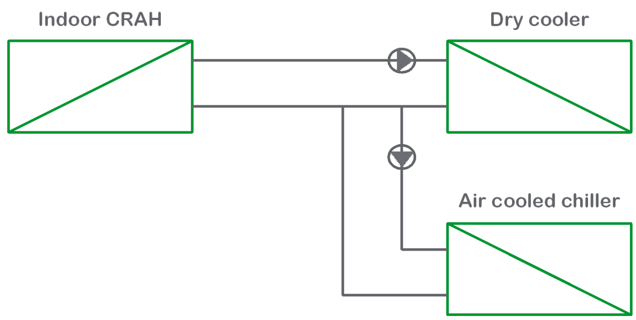

Figure 2

Packaged air-cooled chiller architecture analyzed

The main assumptions we used in the analysis are as follows:

- 1MW data center, fully loaded
- 3 air-cooled chillers[2] in an N+1 configuration, sized for 20 year extreme temperature
- All chillers operate at part load under normal operation (including the redundant chiller)
- Chillers are capable of operating at higher chilled water temperatures. (see **side bar**)
- Variable frequency drive (VFD) dry cooler for economizer mode (no evaporative cooling used)
- Fixed speed pumps
- CRAHs[4] with hot aisle containment in an N configuration
- Airflow demand of servers was matched with CRAH airflow supply (i.e. cfm of servers = cfm of CRAH fans)
- Power density of 4kW/rack
- 3% cost of capital used for TCO calculations
- $0.10 per kilowatt hour cost of electricity
- Weather bin data from ASHRAE Weather Data Viewer 5.0

Operating temperatures of chillers

Every chiller has a maximum chilled water temperature it is capable of supplying. This is limited by the type and design of the chiller. For example, in centrifugal chillers[3] the compressor must be capable of reducing its speed to produce lower refrigerant pressures without damaging the motor or without leaking its lubricating oil into the refrigeration circuit. Depending on the chiller type, other chiller components may require special features which allow for higher chilled water temperatures. We recommend you consult with your chiller vendor before increasing your chilled water set point.

Note that if you have a high efficiency chiller located in a mild climate, the chiller energy savings gained by increasing IT inlet temperatures may not be enough to justify the increase in energy from other devices in the data center (e.g. CRAH, dry cooler, cooling tower, IT).

[1] White paper 132, *Economizer Modes of Data Center Cooling Systems* describes other cooling architectures.

[2] Chiller specifications were from BREC Uniflair chillers because they are specified for data center applications and the data was readily available.

[3] http://www.plantservices.com/articles/2008/258/

[4] CRAH specifications were from InRow CW units because the data was readily available to us. Choosing another air handler product (i.e. room-based or row-based) would not have substantial impact on the results.

We created three different operating temperature scenarios and compared the energy consumption and TCO of each.

- In the **baseline case**, we assumed a fixed IT inlet temperature of 20°C (68°F), which is a typical operating point for data centers today.
- The **second case** allowed temperatures to float from 15.6-26.7°C (60-80°F).
- The **third case** fixed the temperature at 26.7°C (80°F).

We then analyzed the data center in three U.S. cities: Chicago, Seattle, and Miami, to illustrate the impact of varying climates on the results.

Analysis methodology

We analyzed energy cost and capital expense of the entire cooling system utilizing the following methodology:

1. Bin data from ASHRAE Weather Data Viewer 5.0 was used to calculate the cooling system energy at every 1.11°C bin (2°F bin) using thermodynamic principles/formulas. Inputs to this model included CRAH coil effectiveness, average delta T across IT equipment, and equipment losses for IT fans, chillers, dry cooler, pumps, and CRAHs.
2. The 20-year extreme temperature was used as the worst case outdoor temperature for sizing the packaged chiller. This design point is the generally accepted practice for sizing chillers, and recommended by the Uptime Institute.[5]
3. The cooling system energy is dependent on the different operating modes; full mechanical cooling, partial economizer mode, and full economizer mode. The number of hours spent in each operating mode was calculated[6].
4. The IT inlet air set point was used to calculate the chilled water temperature. The chilled water temperature was allowed to range from 7.3°-32°C (45°-90°F).
5. For IT inlet temperatures above 20°C (68°F), the increase in server energy consumption was added to the total cooling system energy consumption. The increase in IT server power consumption was estimated using **Figure 5** and the increase in IT server airflow (CFM) was estimated using the midpoint of the graph in **Figure 6**.
6. The floating temperature scenario represented an ideal (best) case where the chiller and economizer controls allow chilled water temperatures to reset dynamically. In most data centers, the chilled water temperature is set at a fixed temperature year round and would yield lower energy savings than this model projects.
7. The capital expense values were estimated using component, labor, and design prices typically seen in a 1MW data center project. Most of this data came from the *Data Center Capital Cost Calculator*. The change in CRAH capital expense as the IT CFM changes with IT inlet temperature was also accounted for.

Findings

First, we compare the findings of the baseline – where IT temperatures are fixed at 20°C (68°F) – to the second case – where IT temperatures float up and down. Following these

[5] Uptime Institute's "Data Center Site Infrastructure Tier Standard: Topology," http://www.gpxglobal.net/wp-content/uploads/2012/10/TIERSTANDARD_Topology_120801.pdf

[6] In full economizer mode operation, the outdoor conditions allow for all mechanical cooling (i.e. those components used in the refrigeration cycle) to be turned off to conserve energy while still effectively cooling the defined load. When the outdoor temperature limits full economizer mode operation, the cooling plant enters a partial economizer mode of operation, where a proportion of the cooling is handled by the economizer mode and the remaining is handled by the mechanical system. The proportion of each changes (increasing the mechanical cooling proportion as temperature increases outdoors) until full mechanical system operation is required.

findings, we present the comparison of the baseline to the third case – where IT temperatures are fixed at a higher temperature of 26.7°C (80°F).

Baseline vs. floating temperatures
The TCO differences of the baseline vs. the floating temperature case are presented in **Figure 3.** Note – the TCO shown excludes the capital cost of systems that don't change between the two scenarios, i.e. chiller, dry cooler. From this analysis, we can conclude the following:

- While chiller energy always improves (decreases), the net energy consumed does not always improve.
- Higher IT inlet temperatures cause an increase in IT equipment airflow which decreases the deltaT across the CRAHs. More CRAH airflow is needed to remove the same amount of heat at these lower deltaT values.
- The required CRAH capacity increases at higher chilled water supply temperatures because the heat removal capacity of the coil decreases as deltaT decreases.
- The degree to which the increase in energy occurs for the servers and CRAHs depends on the IT equipment characteristics. This is explained in the following section.
- Bin weather data is a significant driver in determining whether floating temperatures from 15.6-26.7°C (60-80°F) results in a cost savings.

Figure 3

Summary of results from baseline of 20 °C (68°F) fixed to floating from 15.6°- 26.7°C (60 °-80 °F) AT FULL LOAD

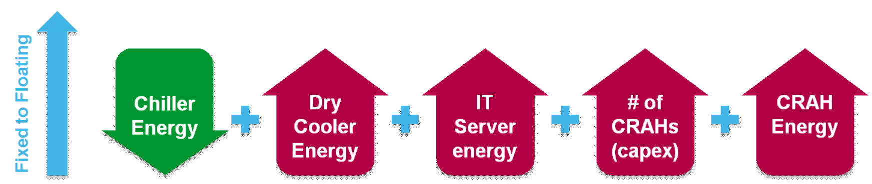

Chicago Fixed	$1,186,000	$162,000	$0	$571,000	$258,000	1% energy reduction, No change in TCO
Chicago Floating	$855,000	$241,000	$173,000	$583,000	$320,000	
Seattle Fixed	$1,196,000	$313,000	$0	$571,000	$258,000	15% energy reduction, 6% better TCO
Seattle Floating	$565,000	$423,000	$175,000	$583,000	$321,000	
Miami Fixed	$2,420,000	$13,000	$0	$571,000	$258,000	11% energy increase, 6% worse TCO
Miami Floating	$2,224,000	$54,000	$342,000	$583,000	$386,000	

Additional results are illustrated in **Table 1**, including differences in the total energy (kWh) and partial power usage effectiveness (pPUE)[7]. **Although PUE improved in all cases, energy did not always improve. This points out the limitation of using *only* PUE as a basis for operational decisions.**

We've also highlighted the maximum float temperature that would result in the lowest TCO for each of the three cities. As the data demonstrates, this optimal temperature varies quite a bit from one city to the next.

[7] In this analysis, pPUE represents only the cooling system losses.

Table 1

Summary of results from baseline of 20˚C (68°F) fixed to floating from 15.6°-26.7°C (60˚-80˚F) AT FULL LOAD

	Chicago	Seattle	Miami
Total energy (kWh)	0.98% savings	14.51% savings	11.01% increase
pPUE (cooling only)	Improves from 1.203 to 1.178	Improves from 1.222 to 1.166	Improves from 1.327 to 1.312
Floating temperature range with lowest TCO	15.6°- 23°C (60° - 74°F)	15.6°- 26.7°C (60˚- 80˚F)	15.6°- 21°C (60°- 70°F)

Figure 4 is a graph illustrating the TCO ($) as we varied the maximum float temperature. In all cases, the *minimum* float temperature was assumed to be 15.6˚C (60˚F).

Figure 4

Summary of results from baseline (68°F) fixed to varied max floating from 15.6-26.7˚C (60 -80°F) AT FULL LOAD

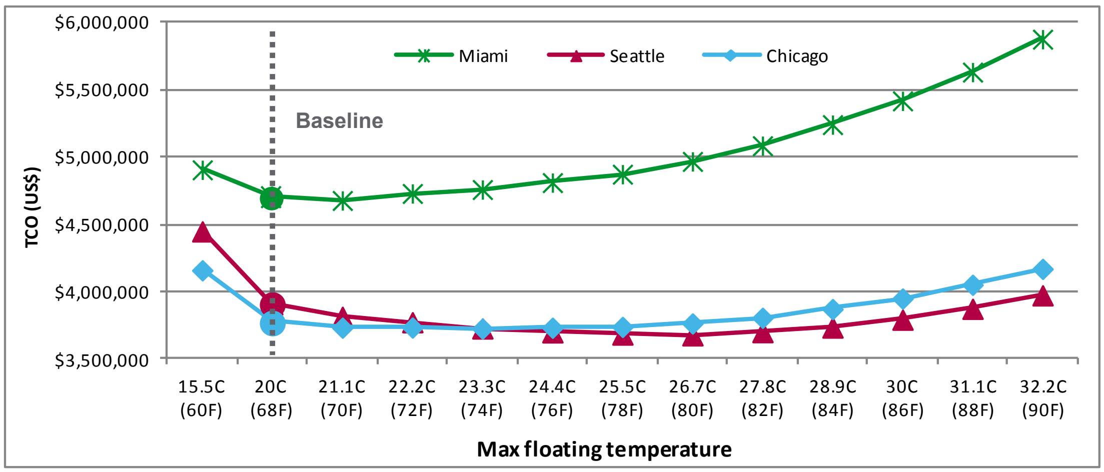

This graph demonstrates how **bin data can have a significant impact on the results**. In Seattle, the optimal temperature for this cooling architecture occurs at 27˚C (80˚F), whereas in Chicago, this same architecture has an optimal maximum temperature of 23˚C (74˚F), and in Miami, that temperature is only 21˚C (70˚F). These findings may come as a surprise to many, but it is driven by the increase in server and CRAH energy that more than offsets the chiller savings above these temperatures. In Miami, the economizer hours are limited by the weather, and so the chiller savings couldn't offset the increase, even at 22˚C (72˚F). **Figure 5** illustrates the bin data of the three cities to show how the distribution of temperatures varies significantly from location to location. The amount of hours at different temperature bins drives how many economizer hours you can gain when you raise the set point.

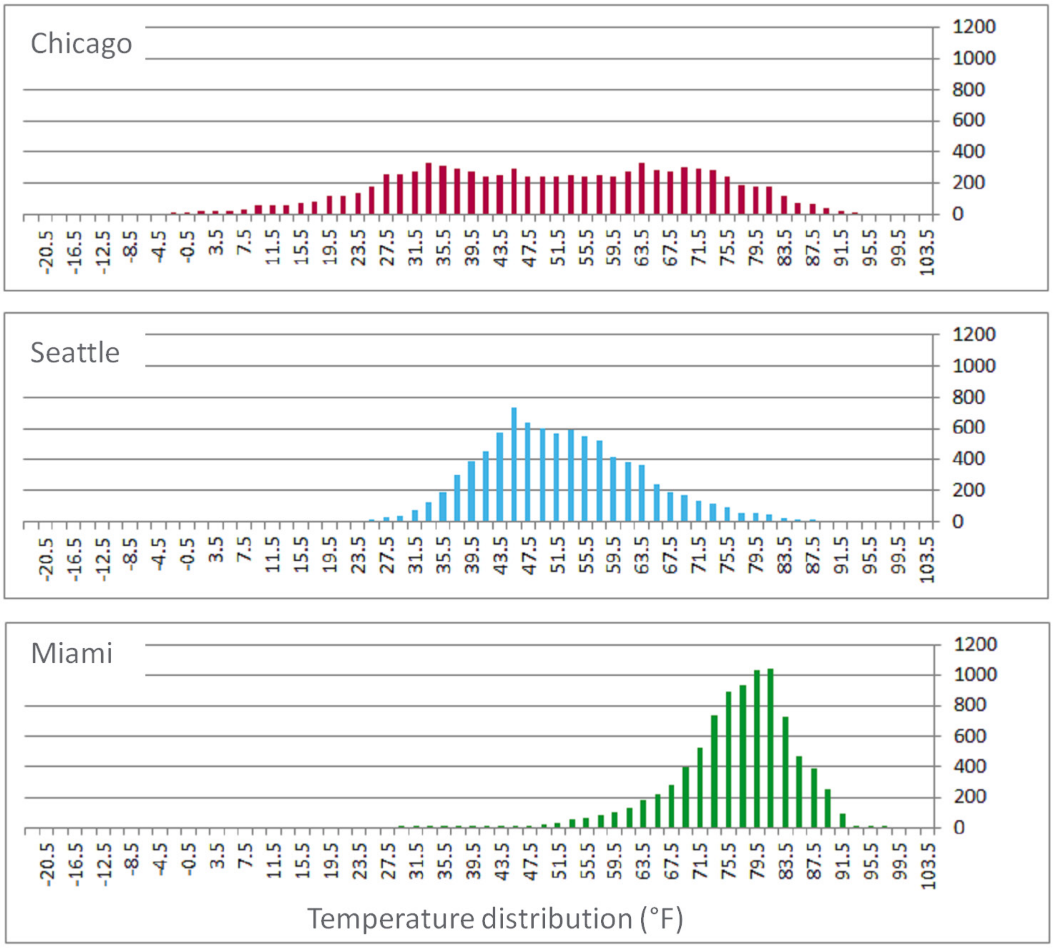

Figure 5

Variation in BIN data for 3 locations: Chicago, Seattle, & Miami (vertical axis represents number of hours)

Baseline vs. higher fixed temperature

When operators think about raising temperatures in their data center, it is commonly understood to mean raising the temperature to a new FIXED set point. Control systems are rarely set up to handle the condition of floating, as the analysis in this paper suggests. So, the question is, what is the impact on energy, TCO, and reliability (X-factor) if the data center temperature were to be raised and FIXED at 27˚C (80˚F)?

The server fans will always draw greater power than the baseline scenario because the higher fixed IT inlet temperature forces the IT fans to spin at the same faster speed all year round. **Figure 6** illustrates how the higher fixed temperature compares to the baseline fixed temperature. The findings are:

- Server energy is even higher than the floating temperature scenario because the server fans are running at the higher temperature year-round.
- Bin weather data is a significant driver in determining whether a higher operating temperature is a smart move.
- Fixing at a higher temperature is always worse than allowing the space to float to that same higher temperature, because when the temperature is fixed, there are never days when the servers and CRAHs can consume less energy.
- No impact on the number of economizer hours and therefore the chiller and dry cooler power consumption, relative to the floating temperature scenario.

Figure 6

Summary of results from baseline of 20˚C (68°F) fixed to a higher fixed temperature of 27˚C (80°F) AT FULL LOAD

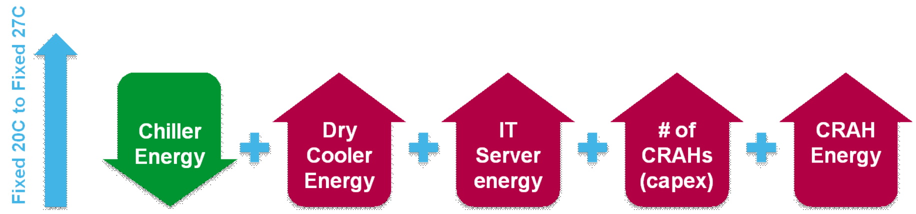

	Chiller Energy	Dry Cooler Energy	IT Server energy	# of CRAHs (capex)	CRAH Energy	
Chicago 20C	$1,186,000	$162,000	$0	$571,000	$258,000	13% energy increase, 5% worse TCO
Chicago 27C	$855,000	$241,000	$345,000	$583,000	$388,000	
Seattle 20C	$1,196,000	$313,000	$0	$571,000	$258,000	2% energy reduction, 1% better TCO
Seattle 27C	$565,000	$423,000	$345,000	$583,000	$388,000	
Miami 20C	$2,420,000	$13,000	$0	$571,000	$258,000	11% energy increase, 6% worse TCO
Miami 27C	$2,224,000	$54,000	$345,000	$583,000	$388,000	

Table 2 summarizes additional findings including total energy (kWh), and pPUE. Again, this illustrates that (1) energy is not always improved when you raise IT temperatures, and (2) PUE as a metric alone is insufficient.

Table 2

Summary of results from baseline of 20˚C (68°F) fixed to FIXED of 27˚C (80°F) AT FULL LOAD

	Chicago	Seattle	Miami
Total energy (kWh)	12.52% increase	2.34% savings	11.16% increase
pPUE (cooling only)	Improves from 1.221 to 1.182	Improves from 1.241 to 1.170	Improves from 1.363 to 1.312

Impact of server fan power & CFM

As the inlet temperature of servers rise, the airflow (CFM) requirement and fan power increases. White Paper 138, *Energy Impact of Increased Server Inlet Temperature* presents the results of a study with energy measurements done in a lab environment of various models of servers. The chart in **Figure 7** is the composite curve from those measurements. Our analysis used this curve as the assumed ramp-up of power draw as temperature increased.

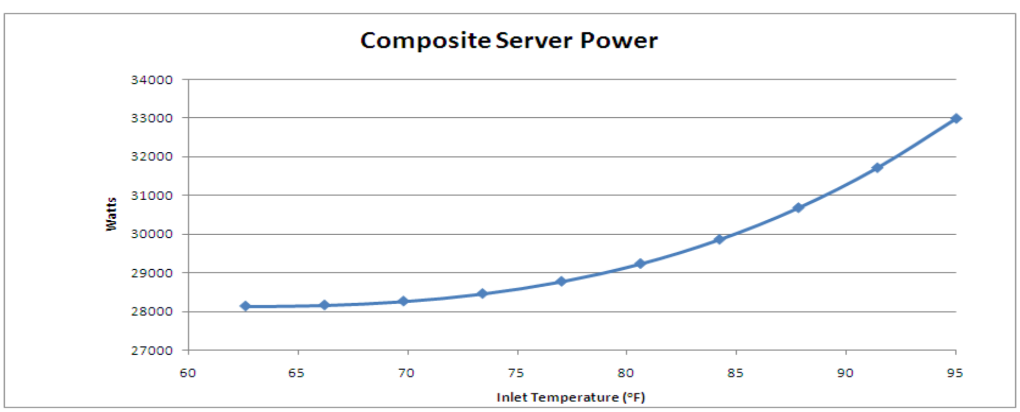

Figure 7

Composite server power vs. inlet temperature as presented in White Paper 138.

The CFM increase as a function of temperature is based on the ASHRAE published data in **Figure 8**. Our analysis used an average curve (dotted red line in the graph), but as the graph illustrates, there is a fairly large variation from server to server as you get to higher temperatures.

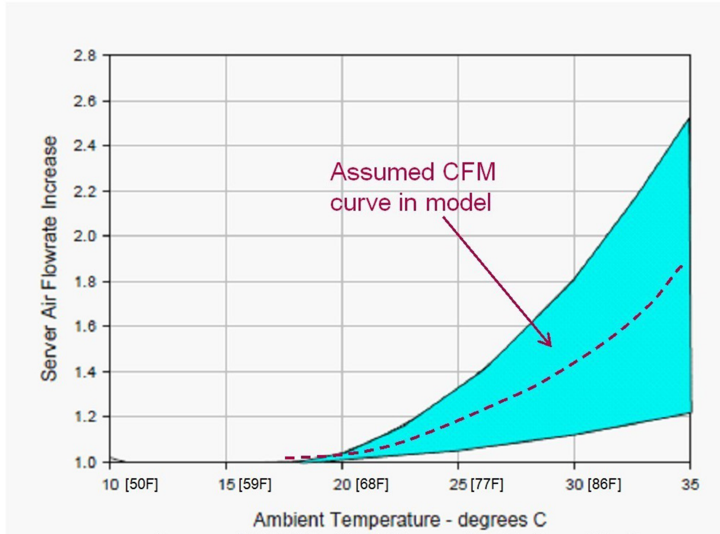

Source: ASHRAE 2011 Thermal Guidelines, Fig 4

Figure 8

Server airflow requirement as a function of operating temperature

If the server CFM requirement didn't ramp-up as temperature increased (meaning if the curve was flat), the results of this analysis would be very different. The IT equipment's behavior at elevated temperatures is what offsets the chiller energy savings, making it a complex analysis. A flat curve would mean higher temperatures are always better because you gain savings through economization with no energy penalty on the CRAH and server side.

To illustrate the impact that the CFM curve has on the overall results, we performed a sensitivity analysis (**Table 3**), with constant chilled water flow and varying the CFM rise as a function of temperature from flat (i.e. no rise) to the highest rise (top curve of the blue area in **Figure 8**). The following occurs as we move to a steeper curve:

- Server fan power becomes a greater penalty because power is proportional to the cube of the shaft speed.
- Number of CRAHs needed increases because you need more airflow.
- CRAH energy increases because you need more airflow.
- Economizer hours go down because you need colder chilled water to make up for the decrease in CRAH deltaT and the associated decrease in CRAH coil effectiveness.

In all three cities, the IT equipment behavior is a key driver to the overall energy impact of going to higher (floating) temperatures. This illustrates the importance of understanding the behavior of your IT gear and analyzing the data center holistically before making operational changes.

Table 3

Impact on total energy (kWh) of varying CFM curves from baseline of 20˚C (68°F) fixed to floating from 15.6°-26.7°C (60˚-80˚F) AT FULL LOAD

		Chicago	Seattle	Miami
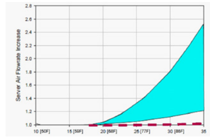	no increase	23% savings	36% savings	13% savings
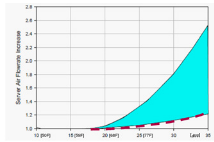	low increase	12% savings	27% savings	1% increase
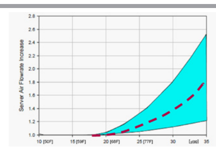	moderate increase	1% savings	15% savings	11% increase
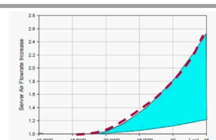	high increase	15% increase	10% increase	25% increase

Impact on reliability

The analysis thus far has focused on the optimal temperature in terms of energy and TCO savings, but reliability is another factor that must be considered when selecting the operating temperature(s). X-factor[8], a metric published by ASHRAE TC9.9 committee, is the ratio of failure rate at a given dry bulb temperature to the failure rate at 20˚C (68˚F). See **Figure 9**.

Figure 9

ASHRAE's X-factor as function of IT inlet temperature

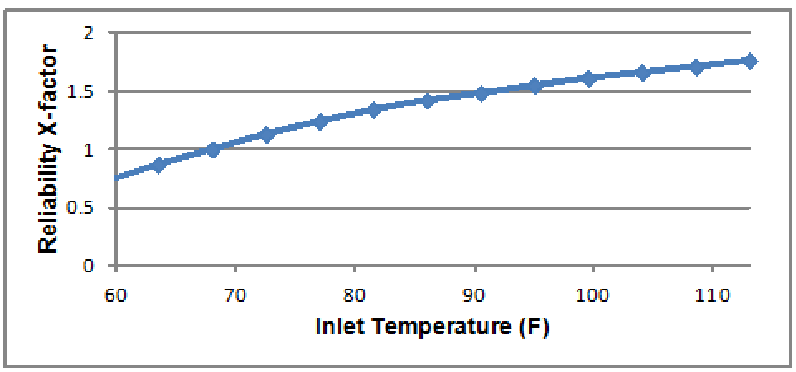

This data illustrates that, relative to the failure rate of servers at 20˚C (68˚F), there will be an increase in failures as the operating temperature rises. Therefore, simply raising a FIXED set point temperature will always decrease reliability if the servers follow the curve of **Figure 7**.

Floating temperatures up AND down is the only way to maintain reliability. As an example, let's say my data center was at 61˚F for half of the year (X-factor=0.8) and 75˚F for the other half of the year (X-factor=1.2), my average X-factor would equal 1. In other words, I'd have no impact on my failures overall.

Is X-factor significant?

Data center operators may question the importance of a change in X-factor given the following:

- Obtaining actual failure rate data from vendors is difficult as it is not generally published data. A 30% X-factor increase in a tiny failure rate number may not be a concern to an operator.
- Failure rates vary by time (i.e. higher failure rate for longer refresh cycles).
- Equipment other than servers (like storage devices) may experience different rates.

[8] http://tc99.ashraetcs.org/documents/ASHRAE%20Networking%20Thermal%20Guidelines.pdf

Figure 10 demonstrates the impact that the maximum float temperature has on X-factor for each of the cities analyzed. This data shows that in Chicago, floating the IT environment up to 23.3˚C (74˚F) enables cost savings without any reliability penalty, and beyond this temperature, there will be an increase in failures relative to the baseline. For Seattle, this temperature is 21.1˚C (70˚F), and for Miami, it is 20˚C (68˚F). This again is driven largely by the bin weather data. If there are a lot of cooler temperature-hours (like in Chicago), they can offset the warmer temperature-hours, but in a more tropical environment (like in Miami), there aren't as many cool temperature-hours to counter those above 20˚C (68˚F).

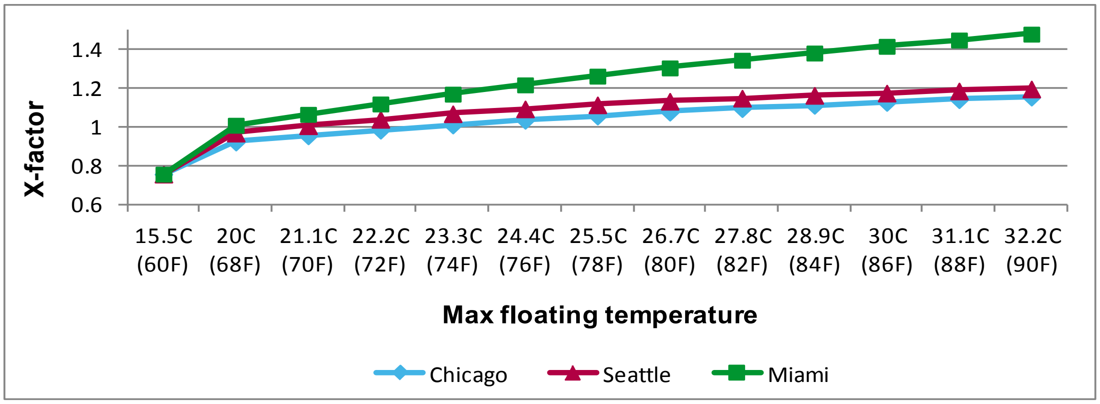

Figure 10

X-factor as a function of floating temperature for Chicago, Seattle, and Miami

When comparing the baseline scenario to the higher fixed temperature of 27˚C (80˚F), there is a **31% increase in failures**. This is regardless of location, because now, the IT equipment is exposed to the same higher temperature year-round.

Another common discussion point when it comes to reliability implications of raising IT temperatures is what happens in the event of a power outage. If my data center is at a higher initial temperature, I have less ride-through time if my cooling system is down before I overheat and crash my IT equipment[9].

Unfortunately, today there seems to be a lack of quantified data on the subject of reliability implication, and while these relative metrics are useful, they are incomplete.

Alternative scenarios

The analysis presented above was based on a particular architecture with particular assumptions. Two key variations are addressed below because they are common occurrences in today's data centers: Oversized CRAHs and lightly loaded data centers.

What if my CRAHs were oversized?

In the analysis described above, we assumed that the CRAH airflow (CFM) was perfectly matched to the IT server airflow requirement, which is the best case from a capital expense perspective. However, this almost never happens in practice because there is always some portion of the cool air that bypasses the IT equipment inlets. In an actual data center, the installed CRAH airflow capacity is always greater than that required by the IT equipment to insure that all IT equipment receives the proper amount of cool air. Some of this oversizing may be intentional, as a safety margin or for redundancy, and some is accidental because of difficulty in forecasting loads or shrinking loads due to virtualization. Uptime Institute

[9] See White Paper 179, *Data Center Temperature Rise During a Cooling System Outage*.

assessments[10] have found this CRAH oversizing to be on average 2.6 times that required by the IT equipment. This oversizing is obviously a capital expense penalty but can actually reduce the energy consumption compared to the ideal "perfectly matched" case.

This is due to the fan laws (sometimes referred to as the cube losses) where fan power is proportional to the cube of the fan shaft speed. When CRAH airflow is oversized, the variable speed fans operate at a lower CFM (i.e. lower speed), therefore consuming less energy. We analyzed the 10-year cooling energy implication of oversizing the CRAH airflow by 25%, 50%, 75%, and 100% while floating the IT inlet temperature up to 27°C (80°F).

Figure 11 shows that while all 3 cities experienced an energy reduction as the oversizing increased, Miami exhibited the largest energy reduction (steeper slope). This is because Miami experienced a limited number of hours at colder temperatures where the fans could reduce their speed. Therefore, a reduction in fan energy by oversizing the CRAH units was realized for nearly all bin hours. Note that this CRAH oversizing comes with an increase in capital expense that typically exceeds the 10-year cooling energy savings. While some oversizing helps prevent hot spots in front of IT equipment, this practice must be balanced with proper air management practices.

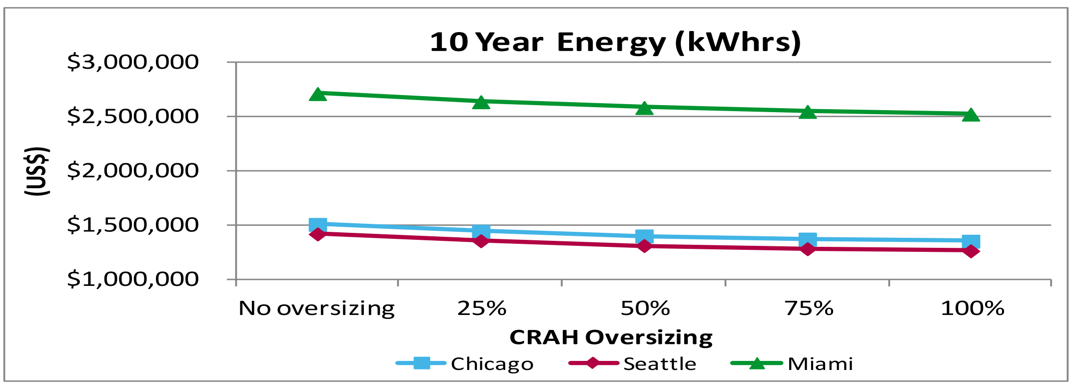

Figure 11

Effect of CRAH oversizing on 10 year energy cost floating from 15.6-26.7˚C (60 -80°F)

What if my data center was only at 50% load?

This is certainly a valid question, as most data center capacities (kW) are based on uncertain future loads, resulting in systems that are under-utilized in practice. The majority of data centers operate between 30 and 60% load.

We ran our same analysis as above, but with the 1 MW data center loaded to 50% (500 kW), and an additional 25% CRAH capacity. The results are illustrated in **Table 4** (baseline vs. floating temperatures) and **Table 5** (baseline vs. fixed higher temperature).

When we float temperatures in a 50% loaded data center, the savings (as a percent of the energy at the fixed baseline temperature) improves. The majority of the additional savings comes from having less chiller energy which is the result of more free cooling hours. This happens because a dry cooler that is half loaded is capable of attaining the chilled water temperature earlier in the year (smaller approach temperature).

It's important to remember that these savings are attainable IF the data center temperature can float. In practice, this is almost never done because control systems are not set up to adjust temperatures dynamically/automatically. **Table 5** shows the results of the 50% loaded data center when you raise the IT space to a fixed temperature of 27˚C (80˚F). Compared to

[10] https://uptimeinstitute.com/uptime_assets/c7f39bad00527fa4e2207a5f1d5dfc1f8295a0a27287bb670ad03fafbdaa0016-00000web4.pdf

the baseline fixed temperature, this represents an energy penalty in all three cities analyzed. In addition, this represents an increase in the X-factor, since the IT equipment is exposed to a constant higher temperature. As mentioned earlier, this can be a reliability concern. In all scenarios, PUE improved which points to the limitation of using this as a sole metric in making decisions.

Table 4

Summary of results from baseline of 20˚C (68°F) fixed to floating from 15.6°-26.7°C (60˚-80˚F) for ***50% LOADED data center***

	Chicago	Seattle	Miami
Total energy (kWh)	12.7% savings	10.9% savings	0.2% increase
pPUE (cooling only)	Improves from 1.092 to 1.075	Improves from 1.091 to 1.079	Improves from 1.157 to 1.138
Temperature with lowest TCO	27°C (80°F)	27°C (80°F)	21°C (70°F)
X-factor	Improves from 1 to 0.94	Improves from 1 to 0.87	Worsens from 1 to 1.27

Table 5

Summary of results from baseline of 20˚C (68°F) fixed to FIXED of 27˚C (80°F) for ***50% LOADED data center***

	Chicago	Seattle	Miami
Total energy (kWh)	5.3% increase	10.6% increase	1.8% increase
pPUE (cooling only)	Improves from 1.092 to 1.075	Improves from 1.091 to 1.080	Improves from 1.157 to 1.138
X-factor	Worsens from 1 to 1.309	Worsens from 1 to 1.309	Worsens from 1 to 1.309

There are several factors that go into the percent improvements/penalties shown in these tables. A 50% loaded data center has an oversized dry-cooler which enables us to get more economizer hours, which means less time on chiller. This impacts not only the dry cooler and chiller energy, but also (in the floating case) the IT kWh penalty. These changes are location specific (BIN data specific), and as these dynamics change, the big drivers to the total energy also change.

For this reason, you may find results in this paper that seem counter-intuitive. Remember that all results shown are *relative* savings/penalties compared to the baseline for the particular location and load.

Recommendations

The analyses of this paper demonstrate that there are many variables that influence cost savings (or penalty), and that raising temperatures is not always a good thing. Before making temperature changes to a data center, it is important to have a solid understanding of the design conditions, system attributes, load, and so on. We recommend the following be done before raising data center temperatures:

- Air management practices such as containment and blanking panels must be in place before attempting to increase IT inlet temperatures. This will avoid creating hot spots. See White Paper 153, *Implementing Hot and Cold Air Containment in Existing Data Centers*, for more information on implementing these practices.
- Make sure you understand how your IT equipment will behave (power consumption and CFM requirement) as you raise temperatures. Ask your IT vendors for this information.

- Consider whether you can adjust the BIOS settings of your IT equipment to optimize their performance at higher temperatures. This requires a higher level of collaboration between facilities and IT departments.
- X-factor predicts a relative increase of failure rates but work with your IT vendor(s) to determine if the actual rate is significant enough to be a concern.
- Since data centers are not solely made up of servers, make sure you also understand the reliability impact on other equipment like storage and networking.
- Ensure your cooling architecture can operate at elevated temperatures (i.e. some chillers cannot run at higher chilled water temperatures).
- Make sure your growth plan comprehends the potential negative energy impact of increasing IT inlet temperatures. In other words, a savings at 50% load might actually be a cost penalty at 80% load.
- Model out how much total energy you may save by raising temperatures vs. other optimization strategies. Companies such as Romonet[11] have software to help analyze the system dynamics of your specific data center. This is critical because every data center will behave differently.
- When evaluating changes, be sure to look at total energy consumption as a metric, as PUE alone can be misleading.

Conclusion

Data center operators are struggling with the decision of raising temperatures in their IT space. Is it safe to do this? What is the right temperature? Is it worth the increased risk? These are some of the questions they are faced with. This paper helps to explain the implications of making the choice to raise IT temperatures. It is important that the architecture be fully understood and that a complete analysis is done before choosing the operating points. This analysis demonstrated that:

- The cooling architecture and geographic location (specifically the temperature profile of the climate) has a significant impact on the optimal IT temperature set point.
- The shape of the server fan and CFM curve are key drivers.
- While raising temperatures improves the chiller efficiency (by increasing economizer hours), that savings can be offset by an increase in IT energy consumption as well as the air handlers.
- Operating conditions like percent load and CRAH oversizing/redundancy influence whether you see a savings or cost penalty.
- You shouldn’t assume that raising the temperature is always a good thing. Understand your specific system dynamics completely before making changes.
- Cooling architectures that use direct and indirect air economizer modes will likely perform better than the packaged chiller architecture we analyzed in this paper.

[11] http://www.romonet.com

About the authors

Wendy Torell is a Senior Research Analyst at Schneider Electric's Data Center Science Center. In this role, she researches best practices in data center design and operation, publishes white papers & articles, and develops TradeOff Tools to help clients optimize the availability, efficiency, and cost of their data center environments. She also consults with clients on availability science approaches and design practices to help them meet their data center performance objectives. She received her bachelor's of Mechanical Engineering degree from Union College in Schenectady, NY and her MBA from University of Rhode Island. Wendy is an ASQ Certified Reliability Engineer.

Kevin Brown is the Vice President of Data Center Global Solution Offer & Strategy at Schneider Electric. Kevin holds a BS in mechanical engineering from Cornell University. Prior to this position at Schneider Electric, Kevin served as Director of Market Development at Airxchange, a manufacturer of energy recovery ventilation products and components in the HVAC industry. Before joining Airxchange, Kevin held numerous senior management roles at Schneider Electric, including Director, Software Development Group.

Victor Avelar is the Director and Senior Research Analyst at Schneider Electric's Data Center Science Center. He is responsible for data center design and operations research, and consults with clients on risk assessment and design practices to optimize the availability and efficiency of their data center environments. Victor holds a bachelor's degree in mechanical engineering from Rensselaer Polytechnic Institute and an MBA from Babson College. He is a member of AFCOM.

Resources

Fundamental Principles of Air Conditioners for Information Technology
White Paper 57

Economizer Modes of Data Center Cooling Systems
White Paper 132

Energy Impact of Increased Server Inlet Temperature
White Paper 138

Cooling Entire Data Centers Using Only Row Cooling
White Paper 139

Data Center Temperature Rise During a Cooling System Outage
White Paper 179

Browse all white papers
whitepapers.apc.com

Cooling Economizer Mode PUE Calculator
TradeOff Tool 11

Browse all TradeOff Tools™
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm

© 2015 Schneider Electric. All rights reserved.