

How to Fix Hot Spots in the Data Center

White Paper 199

Revision 0

by Paul Lin

Executive summary

Data center operators take a variety of actions to eliminate hot spots. Some of these actions result in short term fixes but may come with an energy penalty and some actions may even create more hot spots. Airflow management is an easy and cost effective way to permanently eliminate hot spots while saving energy and can avoid the capital expense of adding more cooling units. This paper describes the root cause of hot spots, recommends methods to identify them, reviews the typical actions taken, and provides the best practices to eliminate them.

Introduction

According to a study done by Uptime Institute, 10% of racks run hotter than recommended reliability guidelines¹. As data center rack power densities increase to an average of 5 kW/rack or greater, we expect that the percentage of racks with hot spots will increase beyond this value.

If hot spots continue for an extended period of time, it can be a serious threat to not only IT equipment reliability and performance, but also to the warranties or maintenance agreements of hardware manufacturers. Therefore, data center operators need to take action as early as possible to avoid these risks.

First we need to clarify what hot spots really are, their root causes, and how to identify them. In general, after finding hot spots, data center operators take a variety of actions to mitigate them. Some examples include, lowering the cooling unit set point, placing perforated tiles in the hot aisle, and placing pedestal fans in front of the problem racks. These typical actions may fix them temporarily, but may also result in an energy penalty or even create more hot spots.

Understanding the root cause of hot spots is critical to fixing them while simultaneously getting other benefits such as saving energy and avoiding the capital expense of adding more cooling units. This paper describes the root cause of hot spots, recommends methods to identify them, reviews the typical actions taken, and provides the best practices to eliminate them.

What are hot spots?

A lot of IT professionals often read temperatures in the hot aisle or the wrong locations in the cold aisle and think they have a hot spot. Then, they take some actions to fix them which may *not* eliminate the hot spots and may even create more of them. Clarifying what hot spots really mean, their root causes, and how to identify them are critical to eliminating them.

Definition

A hot spot is not any random hot temperature inside of a data center. We define a hot spot as a location at the intake of IT equipment where the measured temperature is greater than the expected value as recommended by ASHRAE TC 9.9². Hot spots occur most often towards the top of a rack. The recommended and allowable temperature ranges at the location of the server inlet can be found in ASHRAE's thermal guidelines.

Root causes

Data centers often have excess installed cooling capacity, particularly when the cooling capacity is determined by the rated "nameplate" capacity of the IT equipment. If this is the case, why do hot spots appear? The real reason is inadequate use of cooling capacity and **NOT** because of inadequate cooling capacity or that the heat load is excessive. In other words, there is enough cooling capacity, but not enough of the cooling capacity can get to where it's needed due to the lack of airflow management.

¹ R.F. Sullivan, L. Strong, K.G. Brill, Reducing Bypass Airflow Is Essential for Eliminating Computer Room Hot Spots, The Uptime Institute, Inc., 2004, 2006.

² ASHRAE. 2011, Thermal Guidelines for Data Processing Environments, Developed by ASHRAE Technical Committee 9.9.

Figure 1 shows an example of inadequate use of cooling capacity, which is based on a real case study performed by Schneider Electric. This involves a legacy data center with perimeter cooling. The raised floor and drop ceiling are used as airflow supply and return plenums. First, the cold air is supplied to the floor plenum with a certain pressure and velocity by the CRAH units. Then, the cold air leaves the floor plenum into the IT space through perforated tiles (54% of CRAH airflow) and open cutouts (46% of CRAH airflow) of the raised floor (i.e., leakage air).

The leakage airflow from the cutouts results in a loss of cooling capacity because the air cannot get to the front of the IT equipment and, instead, is bypassing that equipment. In essence, it is short-circuiting back to the cooling units without transferring any heat energy.

Most of the airflow from the perforated tiles (96.29% of perforated airflow) flows through the equipment in the IT racks, but not all of it due to the lack of airflow management. A small part of the cold air (3.71% of perforated airflow) bypasses the IT equipment, short-circuiting back to the cooling units. This bypass air, like the leakage air, results in a loss of cooling capacity. Meanwhile, some “starved” IT equipment is unable to get enough cold air and is forced to make up the difference by drawing in hot exhaust airflow (7.15% of IT airflow) from the back of the rack, which often leads to a hot spot at the front up corner of the “starved” IT equipment. **In summary, practices that reduce leakage, bypass, and recirculation help to eliminate hot spots.**

The laws of energy conservation and mass (airflow) conservation, allow us to calculate the percentages of the different types of airflows described above in this data center:

- 52% of the total CRAH airflow is supplied to the front of IT racks.
- 46% of available cooling capacity is lost due to leakage through cable cutouts, holes under racks, or misplaced perforated tiles, and 2% of CRAH cooling capacity is lost due to bypass air. All these issues are caused by lack of airflow management (i.e. blanking panels, air containment).
- 7.2% of the total IT airflow is recirculated (i.e. the hot IT exhaust airflow recirculates back to the front of the racks, which is one of the main reasons for hot spots).
- 92.9% of the total IT airflow which is 52% of the total CRAH airflow is returned back to the CRAH units through the drop ceiling.

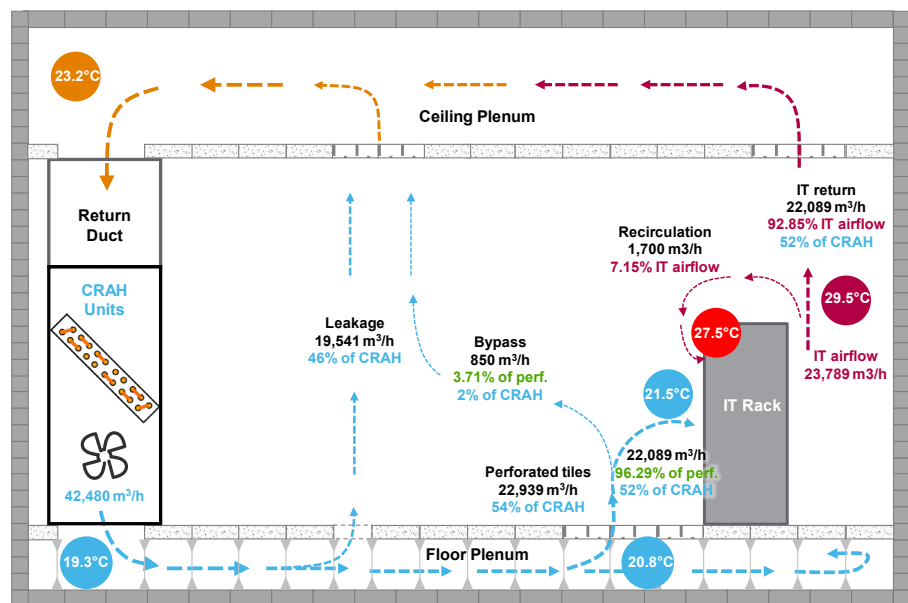


Figure 1

An example of cooling capacity usage assessment in a typical and real data center

How to identify a hot spot

Finding a hot spot as early as possible is critical to prevent IT equipment from overheating and malfunctioning. There are three main ways to detect a hot spot:

- Feel it by walking around
- Manual measurements
- Automatic monitoring

The method that is the lowest in cost and is the easiest for detecting a hot spot is to feel for high temperatures at the front of racks with your hands by walking around. This is also the least accurate method, but is somewhat effective for extreme hot spots.

Manual measurements are a better method since meters provide accuracy and quantify the temperature. Some examples of meters include plastic temperature strips, a temperature gun, and forward looking infrared (FLIR). Manual measurements are regarded as low cost, yet effective ways to locate a hot spot. Even infrared cameras are available today for about \$300 US. Data center operators can use these meters to gather temperature measurements at the server intakes and rack front doors, as well as the temperature difference between server intakes and exhaust vents (i.e. server deltaT) in order to identify a hot spot.

Automatic monitoring is the best method since it can show live data, illustrating the thermal conditions of the server or data center. Automatic monitoring with a data center infrastructure management (DCIM) solution can send real-time alerts via email or text to specific personnel if a threshold is reached. Through DCIM software, you can get actual inlet and outlet temperatures at the device level depending on your requirement. Schneider Electric's StruxureWare™ is an example of this type of DCIM software, which can provide detailed 3D thermal mapping using real-time data from the installed network of sensors. This is the most accurate method with the highest cost. Furthermore, IT devices normally have thermal sensors within them, which can monitor their thermal status and report a high temperature using IPMI protocol.

Some other tips for identifying or preventing *potential* hot spots include:

- Use metered rack PDUs to identify and inspect higher density racks (i.e. greater than 5kW) which are more likely to experience hot spots.
- CFD software to predict hot spots after proposed moves, adds, and changes or during the data center design phase. CFD simulation can provide a detailed 3-D analysis of temperature and pressure contour in front of the rack, and also the airflow distribution around the rack to identify a potential hot spot. The tool's real power is in identifying places where cold air is being wasted or is mixing with hot air, causing inadequate use of the cooling capacity.

Typical actions taken

Data center operators take a variety of actions to fix hot spots after they identify them. However, not all of these actions are effective. The following section describes some typical actions taken and highlight why they are good or ineffective practices. Note that most of these actions do nothing to reduce bypass or recirculation.

- Lower cooling unit set point (Temporary emergency use only)
- Place perforated tiles in hot aisle (Ineffective practice)
- Place racks and perforated tiles close to cooling units (Ineffective practice)
- Place pedestal fans in front of problem rack (Temporary emergency use only)
- Blow air across ice into cold aisle (Ineffective practice)
- Roll in portable cooling unit (Temporary emergency use only)

- Add more cooling units (Likely not required)
- Relocate problem loads (Good practice)

Lower cooling unit set point (Temporary emergency use only)

It seems logical that a lower supply air temperature helps reduce hot spots. However, this should be the last resort to fixing hot spots because it can lower the overall cooling system efficiency and capacity. The effect of this action depends on the percent load of CRAC/CRAH units. If the cooling system has extra capacity (i.e. operating at less than 100% load and not at its thermodynamic limit), then lowering the set point can have a positive effect. When the CRAC/CRAH is in close proximity to the hot spot, the lower set point may reduce the temperature at the hot spot.

If the CRAC/CRAH is already operating at its maximum capacity (100% load), then it is not possible to lower the temperature set point, as the system is already at its thermodynamic limit and will do nothing to fix hot spots. Every cooling system has an inherent maximum capacity for a given set of environmental conditions. This “maximum” capacity is reduced when the set point is reduced. For more information on this topic, please see **Appendix: [The physics behind cooling capacity decrease](#)**.

Place perforated tiles in hot aisle (Ineffective practice)

Some think this is a good practice because they don't understand the benefit of a hot aisle / cold aisle layout and see all high temperatures as a problem. However, this action does not eliminate hot spots in cold aisles, and may even create more of them. Furthermore, placing perforated tiles in the hot aisle (i.e. bypass air) decreases the available cooling capacity. A hot aisle / cold aisle row layout is a best practice and therefore, there is no such thing as a hot spot in the hot aisle. What's important for IT equipment is that the cold aisle remains cold, because it is the “cold air reservoir” from which IT equipment draws upon.

In the early days of air-cooled mainframes, cooling was accomplished using raised-floor cooling with cooling units controlled off of return air temperature. This approach worked due to a uniform, well-mixed room air temperature. Today, hot aisle / cold aisle layouts purposely create areas of hot and cold temperatures which, by design, result in non-uniform return air temperatures. People accustomed to a uniform room air temperature may place perforated tiles in the hot aisle thinking they are solving a hot spot problem.

Place racks and perforated tiles close to cooling units (Ineffective practice)

Some think this is a good practice because they assume racks and perforated tiles placed within a few feet of cooling units receive the most cooling capacity. In fact, this practice may have the opposite effect. It may starve IT equipment of cold air and may not fix hot spots in a consistent way. While this practice may help capture most of the hot exhaust air, this is not predictable, and there are more effective ways of fixing hot spots. The reason why this practice may starve IT equipment is that the air exiting the cooling unit fan is typically moving at a very high velocity which creates low static pressure in that area. This means that the perforated tiles in this area supply very little, if any, cold air and may even suck air down into the subfloor plenum. In general, racks should be deployed at least 3 tiles away from the cooling units.

An easy way to identify whether there is an air velocity problem under the raised floor is to place a small piece of paper on the perforated tile. If the paper is pulled in by the perforated tile, then the tile should be replaced with solid tiles to balance the pressure under the raised floor.

Place pedestal fans in front of problem rack (Temporary emergency use only)

Some think this is a good practice because it focuses airflow directly in front of a specific hot spot. However, this practice should only be used temporarily in emergency situations such as when IT equipment is about to enter thermal shutdown. This action (shown in **Figure 2**) can address a hot spot by decreasing the equipment operating temperature, but at significant cost. Pedestal fans essentially work by operating as a mixer, blending the equipment exhaust air and the cold supply air to create an air temperature that is hotter than the supply, but cooler than the equipment exhaust. This will also act to increase the flow of air through the equipment. The mixing between hot and cold air will decrease cooling system efficiency, leading to increased dehumidification / humidification, inadequate use of cooling system capacity, and possible loss of cooling redundancy. Furthermore, pedestal fans are an additional source of heat in the data center.

Figure 2

Example of pedestal fans placed in front of problem rack to eliminate hot spots



Blow air across ice into cold aisle (Ineffective practice)

Some think this is a good practice because ice is an easy way to store thermal energy. Though this action (shown in **Figure 3**) can help mitigate hot spots it is not worth the risk of spilling containers filled with water as the ice melts. Even if contained ice-packs are used, there are still much more effective and easier solutions, as discussed in the next section.

Figure 3a (left)

Example of fans used to blow air cross ice into cold aisle



Figure 3b (right)

Example of ice boxes stored in a freezer



Roll in portable cooling unit (Temporary emergency use only)

Some think this is a good practice because it focuses cold air flow directly in front of a specific hot spot. However, this practice should only be used temporarily in emergency situations such as when IT equipment is about to enter thermal shutdown. Unfortunately, this practice is often used as a permanent solution. Portable cooling units are generally used in case of a loss of cooling event since it can be rolled into place easily by data center operators. However, the best practices discussed in the next section are the preferred, permanent, and cost effective ways to fix widespread hot spots.

Add more cooling units (Likely not required)

Some think this is a good practice because they automatically associate hot spots with a lack of cooling capacity. However, in most cases, there is enough cooling capacity, but not enough of the cooling capacity can get to where it's needed due to the lack of airflow management. Furthermore, this action does not always fix hot spots but comes with a significant capital investment. According to a study done by Uptime Institute³, despite some IT rooms having 15 times more cooling capacity than required, 7% to 20% of the racks in the room still had hot spots. The reason for this is that the cold supply air is bypassing the air intakes of the computer equipment. The right approach is to adopt the best practices discussed in the next section and then verify if more cooling units are really needed.

Relocate problem loads (Good practice)

This action involves moving IT equipment from one rack to another rack in order to eliminate the hot spot. This is actually a good practice and is discussed in the next section.

Best practices for fixing hot spots

Most of the common practices above are not recommended because they do not address the two main causes of hot spots: bypass and recirculation. Completely eliminating bypass and recirculation means that the hot and cold air streams are entirely separated which means no hot spots can exist. The first four best practices below are effective because they address the issue of bypass air, recirculation, or both. The last best practice should only be done once adequate air management practices are implemented.

1. Manage airflow in the rack
2. Manage airflow in the room
3. Relocate problem loads
4. Change the location of the air temperature sensor
5. Allow DCIM to control airflow of cooling units

Manage airflow in the rack

Many hot spots appear due to hot exhaust air recirculation inside or around the rack. Therefore improving rack airflow management plays an important role in fixing hot spots. Open rack "U" spaces and cable penetrations are a common cause of hot exhaust air recirculation, which results in hot spots. Using blanking panels to fill the unused "U" positions and brush strip for cable penetrations in racks or enclosures is one of the most simple and cost effective methods for improving airflow in the rack. Data center operation procedures should be updated to require that blanking panels and brush strips are installed after all moves, adds, and changes. For more information on airflow management with blanking panel, please see White Paper 44, [Improving Rack Cooling Performance Using Airflow Management Blanking Panels](#).

Some types of switches and routers require side-to-side airflow. When they are deployed in a data center with traditional front-to-back rack airflow, the hot exhaust air from the switch/router may return to its intake and cause hot spots. Rack side air distribution units shown in **Figure 4a** can be used to deliver cool air directly and predictably to equipment with side-to-side airflow without creating hot spots. For more information on this topic, see White Paper 50, [Cooling Options for Rack Equipment with Side-to-Side Airflow](#).

³ R.F. Sullivan, L. Strong, K.G. Brill, Reducing Bypass Airflow Is Essential for Eliminating Computer Room Hot Spots, The Uptime Institute, Inc., 2004, 2006.

Where the overall average cooling capacity is adequate but hot spots have been created in racks with above-average rack density, cooling can be improved by retrofitting fan-assisted devices that improve airflow and increase cooling capacity. Devices shown in **Figure 4b** can effectively “borrow” the air from adjacent racks to support loads up to 3 kW per rack. It can minimize temperature differences between the top and bottom of the enclosure, while also preventing hot exhaust air recirculation to the inlet of the enclosure. As with all air-scavenging devices, care must be taken when positioning the device to ensure that the air taken from the adjacent space does not result in overheating of neighboring racks. These devices should be UPS-powered to avoid thermal shutdown of equipment during power outages. In high density environments, thermal overload can occur during the time it takes to start the backup generator.

Figure 4a (left)

Example of a rack airflow management using the rack side air distribution unit





Figure 4b (right)

Example of a rack airflow management using the rack mounted fully ducted air supply unit



Manage airflow in the room

After improving airflow management in the rack, the next important step is to improve airflow management in the room. To begin, seal all openings in the raised floor. Brush strips should be used to seal cable cutouts behind racks and under PDUs. These openings cause the majority of unwanted air leakages. Seal gaps around cooling units and other floor penetrations with air dam foam or pillows, replace missing floor tiles with solid tiles, and identify perforated tiles that are a source of bypass air and replace them with solid tiles. For example, if the tile is placed in front of an empty rack, replace it with a solid tile. Also, rebalance underfloor airflow by using the procedure in the side bar. Proper tile placement coupled with sealed gaps in the raised floor will reclaim lost cooling capacity.

Another contributor of hot spots is the hot and cold air mixing that occurs over the tops of racks and around the end of racks. A key best practice is to separate the hot and cold air streams by containing the aisle and/or the rack. **Figure 5** shows examples of hot and cold aisle containment applications. Containment not only helps to fix hot spots, but also provides the added benefit of energy savings over traditional uncontained data center designs. For more information on containment deployment in existing data centers, see White Paper 153, [Implementing Hot and Cold Air Containment in Existing Data Centers](#). For more information on containment deployment for a new data center, see White Paper 135, [Impact of Hot and Cold Aisle Containment on Data Center Temperature and efficiency](#).

Rebalancing underfloor airflow

Some think that placing as many perforated tiles as possible in each row can ensure enough cooling capacity. However, too many tiles disrupt underfloor airflow uniformity through the whole IT space. Adding more tiles won't necessarily fix a hot spot, but taking some away from another area will help the air get to the hot racks. The following guidance shows how.

First, get a flow meter, and measure the total CFM in the row. Then, divide by 125 CFM/kW and this should be approximately equal to the rack power in the row. If not, remove some tiles and repeat the above procedures until they are approximately equal. The following picture shows an example of a flow meter.



Figure 5a (left)

Example of a cold aisle containment solution



Figure 5b (right)

Example of a hot aisle containment solution



Figure 6a shows an example of an active individually-ducted rack. The rear door of the rack can be replaced with an air-moving device to convert it to an active ducted rack. Note that these devices add approximately 250 mm (10 inches) to the overall depth of the rack which will likely add to the pitch⁴ between adjacent rows of racks. Hot exhaust air that would normally be expelled into the hot aisle is gathered and moved upwards, where it is ducted into the return air plenum. This eliminates recirculation at the rack and improves cooling system efficiency and capacity. The fans in the active ducting system allow for rack densities up to about 12 kW and can overcome adverse plenum pressure or pressure drops resulting from very dense cabling at the server exhausts. However, active systems can easily create unintended consequences elsewhere in the data center and should be deployed with particular care. Blanking panels and rack side panels must be used with these devices. Active systems consume power and require monitoring and maintenance. For more information on this topic, please see White Paper 182, [The Use of Ceiling-Ducted Air Containment in Data Centers](#).

For high density racks, the equipment located at the top half of the rack normally can't get enough cold air. **Figure 6b** shows an example of a floor-mounted device which can adjust its fan speed to supply a suitable airflow rate depending on rack exhaust air temperature or under-floor pressure. It can be used to support up to 12 kW per rack.

Figure 6a (left)

Example of ducted rack containment using an active individually-ducted rack





Figure 6b (right)

Example of an active tile



Relocate problem loads

As discussed above, relocating problem loads addresses hot spots by identifying the “problem” load and moving it to a lower-density rack. Provide the room with the capability to power and cool to an average value below the peak potential enclosure value, and spread out the load of any proposed enclosures whose load exceeds the design average value by splitting the equipment among multiple rack enclosures. Note that spreading out equipment among multiple racks leaves a sizable amount of unused vertical space within the racks. This space must be filled with blanking panels to prevent degradation of cooling performance. If it is possible to take down a server or other critical piece of equipment, this practice can fix a hot spot with little expense. For more information on this topic, see White Paper 46, [Cooling Strategies for Ultra-High Density Racks and Blade Servers](#).

⁴ The measurement of row-to-row spacing is called pitch, the same term that is used to describe the repeating center-to-center spacing of such things as screw threads, sound waves, or studs in a wall.

Change the location of the air temperature sensor

In most older data centers, thermostats are installed in the CRAC return air stream, where airflow may be unpredictable. This can also result in uneven CRAC unit loading, which, in turn, leads to variations in server inlet temperatures. Relocating the thermostats to the supply air stream, where discharge air can be predictably controlled, will provide more uniform IT equipment inlet temperatures⁵. When combined with containment methods, relocating the thermostat also enables supply air temperatures to be increased thereby saving cooling system energy without the concern for wide variations in supply air temperatures.

Allow DCIM to control airflow of cooling units

Systems are available that control individual room-based cooling units based on the temperatures at the front of IT racks. These systems use learning algorithms to dynamically predict and adjust cooling unit fan speeds and which cooling units can be turned off. Controlling the amount of air into the data center can limit the amount of bypass air. An example of such a system is the [Vigilent](#) cooling system.

Conclusion

Hot spots can result in poor server reliability, performance, or even server damage. Hot spots typically appear at the intake of IT equipment due to inefficient air flow management like cold air leakage (i.e. bypass air) and hot exhaust air recirculation. Feeling it by walking around, taking manual measurements, or having automatic monitoring are three main ways to detect them.

Many actions taken by data center operators to fix hot spots are either temporary emergency use only, ineffective, likely not required, or make matters worse. Rack and room airflow management, containment, relocating problem equipment, changing the location of the air temperature sensors, and allowing DCIM to control airflow of cooling units are the best practices which can help fix hot spots in an easier and cost-effective manner.



About the author

Paul Lin is a Senior Research Analyst at Schneider Electric's Data Center Science Center. He is responsible for data center design and operation research, and consults with clients on risk assessment and design practices to optimize the availability and efficiency of their data center environment. Before joining Schneider Electric, Paul worked as the R&D Project Leader in LG Electronics for several years. He is now designated as a "Data Center Certified Associate", an internationally recognized validation of the knowledge and skills required for a data center professional. He is also a registered HVAC professional engineer. Paul holds a master's degree in mechanical engineering from Jilin University with a background in HVAC and Thermodynamic Engineering.

⁵ <http://www.futurefacilities.com/media/whitepapers/SupplyControl/SupplyControl.pdf>



 [Improving Rack Cooling Performance Using Airflow Management Blanking Panels](#)
White Paper 44

 [Cooling Options for Rack Equipment with Side-to-Side Airflow](#)
White Paper 50


 [Implementing Hot and Cold Aisle Containment in Existing Data Centers](#)
White Paper 153

 [Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency](#)
White Paper 135

 [The Use of Ceiling-Ducted Air Containment in Data Centers](#)
White Paper 182

 [Fundamental Principals of Air Conditioners for Information Technology](#)
White Paper 57

 [Browse all white papers](#)
whitepapers.apc.com

 [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

Appendix: The physics behind cooling capacity decrease

To easily understand why cooling capacity is reduced after lowering the cooling unit set point, consider the purpose of the entire cooling system. The purpose of a cooling system is to collect and transport unwanted heat energy from inside the data center to the outdoors. This can be accomplished using a vapor compression cycle such as direct expansion (DX) or using an economizer mode. According to the Second Law of Thermodynamics⁶, input power must be used to move heat from low temperature environments to high temperature environments. In any case, the key parameter is the difference between the outdoor temperature and the IT supply air temperature. The bigger this temperature difference, the more work the cooling system must do, which reduces the overall capacity of the cooling system (assuming a constant outdoor temperature). **Figure A1** shows the heat transfer from the data center to the outdoors when the temperature difference is small and large.

Figure A1

Heat transfers from data center to outdoors according to the second law of thermodynamics



Consider a chilled water plant. Decreasing the supply air temperature set point increases the difference in temperature between the supply air and the outdoor air. The system must now work harder to reject the same amount of heat from the data center. Similarly, if the set point is kept constant and the outdoor temperature increases, the capacity will decrease⁷, since the indoor / outdoor temperature difference increased. Further evidence to support this fact can be seen in technical specification documents for cooling systems which demonstrate a decrease in maximum capacity at lower set points.

Cooling capacity also decreases for indirect economizer mode systems, again due to the increased indoor / outdoor temperature difference. In other words, the maximum capacity of the exchanger decreases due to the larger temperature difference between the return air (or water) and outdoor air (or water). For systems with direct air economizer modes, the capacity of the system is decreased in the form of lower economizer hours. For more information on cooling fundamentals, please see White Paper 57, [Fundamental Principles of Air Conditioners for Information Technology](#).

⁶ A cyclic transformation whose only final result is to transform heat extracted from a source which is at the same temperature throughout into work is impossible. - Scottish physicist William Thompson (Lord Kelvin)

⁷ [Impact of elevated ambient temperatures on capacity and energy input to a vapor compression system – Literature review](#), NIST, Table 1, page 5. (Accessed May 21, 2014)