Data Center Cooling Best Practices

By Peter Sacco

Experts for Your Always Available Data Center

White Paper #2

EXECUTIVE SUMMARY

Maintaining a suitable environment for information technologies is arguably the number one problem facing data center and computer room managers today. Dramatic and unpredictable critical load growth has levied a heavy burden on the cooling infrastructure of these facilities making intelligent, efficient design crucial to maintaining an always available data center. The purpose of this white paper is to establish a best practices guideline for cooling systems design for data centers, computer rooms, and other mission critical technical spaces.

Table of Contents

Cooling Systems Design Goals

To establish an effective cooling solution for any new or upgraded data center or computer room, it is essential to establish a set of design goals. Experience suggests these goals can be categorized as follows:

Adaptability

1. Plan for increasing critical load power densities
2. Utilize standard, modular cooling system components to speed changes
3. Allow for increasing cooling capacity without load impact
4. Provide for cooling distribution improvements without load impact

Availability

1. Minimize the possibility for human error by using modular components
2. Provide as much cooling system redundancy as budget will allow
3. Eliminate air mixing by providing supply (cold air) and return (hot air) separation to maximize cooling efficiency
4. Eliminate bypass air flow to maximize effective cooling capacity
5. Minimize the possibility of fluid leaks within the computer room area as well as deploy a detection system
6. Minimize vertical temperature gradients at the inlet of critical equipment
7. Control humidity to avoid static electricity build up and mold growth

Maintainability

1. Deploy the simplest effective solution to minimize the technical expertise needed to assess, operate, and service the system
2. Utilize standard, modular cooling system components to improve serviceability
3. Assure system can be serviced under a single service contract

Manageability

1. Provide accurate and concise cooling performance data in the format of the overall management platform
2. Provide local and remote system monitoring access capabilities

Cost

1. Optimize capital investment by matching the cooling requirements with the installed redundant capacity and plan for scalability
2. Simplify the ease of deployment to reduce unrecoverable labor costs
3. Utilize standard, modular, cooling system components to lower service contract costs
4. Provide redundant cooling capacity and air distribution in the smallest feasible footprint

Determine the Critical Load and Heat Load

Determining the critical heat load starts with the identification of the equipment to be deployed within the space. However, this is only part of the entire heat load of the environment. Additionally, the lighting, people, and heat conducted from the surrounding spaces will also contribute to the overall heat load. As a very general rule-of-thumb, consider no less than 1-ton (12,000 BTU/Hr / 3,516 watts) per 400 square-feet of IT equipment floor space.

The equipment heat load can be obtained by identifying the current requirements for each piece of equipment and multiplying it by the operating voltage (for all single phase equipment). The number derived is the maximum draw or nameplate rating of the equipment. In reality, the equipment will only draw between 40% and 60% of its nameplate rating in a steady-state operating condition. For this reason, solely utilizing the nameplate rating will yield an over inflated load requirement. Designing the cooling system to these parameters will be cost prohibitive. An effort is underway for manufacturers to provide typical load rating of all pieces of equipment to simplify power and cooling design.

Often, the equipment that will occupy a space has not been determined prior to the commencement of cooling systems design. In this case, the experience of the designer is vital. PTS maintains an expert knowledge of the typical load profile for various application and equipment deployments. For this reason, as well as consideration of future growth factors it may be easier to define the load in terms of an anticipated standard for a given area. The old standard used to be a watts-per-square foot definition. However, that method has proven to be too vague to be effective.

Establish Power Requirements on a per RLU Basis

Power density is best defined in terms of rack or cabinet foot print area since all manufacturers produce cabinets of generally the same size. This area can be described as a rack location unit (RLU), to borrow Rob Snevely's, of Sun Microsystems, description.

The standard RLU width is usually based on a twenty-four (24) inch standard. The depth can vary between thirty-five (35) and forty-two (42) inches. Additionally, the height can vary between 42U and 47U of rack space, which equates to a height of approximately seventy-nine (79) and eighty-nine (89) inches, respectively.

A definite trend is that RLU power densities have increased every year.

- 1,500-watts per RLU for a typical late-90's server deployment
- 4,000-watts per RLU for a typical easrly-2000's server deployment
- 5,000-8,000-watts per RLU for a 1U server deployment
- Upwards of 30,000-watts per RLU for a blade server deployment

In 2002, American Power Conversion (APC) published data that approximately 90% of all new product server environments were being deployed at rack densities between 1,500 and 4,000 watts.

The reality is that a computer room usually deploys a mix of varying RLU power densities throughout its overall area.

The trick is to provide predictable cooling for these varying RLU densities by using the average RLU density as a basis of the design while at the same time providing adequate room cooling for the peak RLU and non-RLU loads.

Determine the CFM Requirements for each RLU

Effective cooling is accomplished by providing both the proper temperature and an adequate quantity of air to the load.

As temperature goes, the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) standard is to deliver air between the temperatures of 68 °F and 75 °F to the inlet of the IT infrastructure. Although electronics performs better at colder temperatures it is not wise to deliver lower air temperatures due to the threat of reaching the condensate point on equipment surfaces.

Regarding air volume, a load component requires 160 cubic feet per minute (CFM) per 1 kW of electrical load. Therefore, a 5,000-watt 1U server cabinet requires 800 CFM.

Most perforated raised floor tiles are available in 25%-open and 56%-open versions. Typically, 25%-open tiles should be used predominantly and 56%-open tiles used sparingly and for specific instances. Additionally, damped, adjustable-flow tiles are also available, but are not recommended due to the complexity involved in balancing variable air delivery across a raised floor.

Typical raised floor cooling capacities can be determined per the following table:

~300 CFM	1,875 watts	Can be achieved with a standard raised floor cooling design
~700 CFM	4,375 watts	Can be achieved with properly laid out, low leakage, raised floor design
~1,200 CFM	7,500 watts	Can only be achieved within specific areas of a well planned raised floor

Table 1

Raised floor cooling effectiveness can be further enhanced by ducting IT-equipment return air directly to the CRAC equipment such that uniform air delivery across the raised floor is not critical.

Perform Computational Fluid Dynamic (CFD) Modeling

CFD modeling can be performed for the under floor air area as well as the area above the floor. CFD modeling the airflow in a computer room provides information to make informed decisions about where to place CRAC equipment, IT-equipment, perforated tiles, high density RLUs, etc. Much of the software available today also allows mapping of both under floor and overhead airflow obstructions to more accurately represent the environment.

Airflow modeling gives engineers the luxury to consider several design options in the minimum amount of time. As a result, the final design is not based on a tentative approach, but is a result of a professional design process considering several options and selecting the optimum solution. This can save on capital and running costs, save time on correcting mistakes further down the design route, and save time on commissioning.

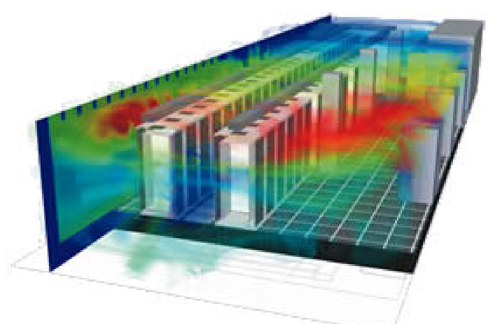

Determine the Room Power Distribution Strategy

The two (2) main decisions in developing a room power distribution strategy are:

1. Where to place the power distribution units (PDUs)
2. Whether to run power cables overhead or under the floor

Traditional design practices dictated placing PDUs on the perimeter of the computer room space or in mechanical rows down the middle of larger rooms. Both approaches include running individually routed branch circuit cables, under the raised floor, from the PDU to the RLUs. In smaller computer rooms PDUs may not even be utilized choosing instead to use wall-mounted power panels. Today however, it has become ever more commonplace to provide power distribution directly at the row of RLUs, or for a select ‘pod’ of RLUs, herein referred to as the, ‘row/pod’ approach.

In the homerun approach, a mass of branch circuit cable can cause cable dams under the floor. This congestion reduces the amount of clear space in the floor plenum. The less open space the air has to flow through, the faster the air travels - The faster the air travels, the less static pressure available - The less the static pressure, the less cold air that is available for the load.

The row/pod PDU approach has the advantage of reducing the number of cables traversing the sub-floors. The affect is greatly reduced clutter under the floor, thereby improving cooling performance. This approach also has the added benefit of the ability to route 480-volt feeder cables to the row/pod PDUs, thereby reducing the diameter of the cables.

Power Cable Pathway Considerations

Overhead power cable pathways and under floor power cable pathways both have an impact on cooling performance. Overhead cable pathways, which are typically located in the hot aisles can be disruptive to the hot air return path to the CRAC units and be just as detrimental to overall cooling effectiveness. Overhead strategies should only be employed when there is ample space between the tops of the RLUs and the underside of the suspended ceiling or deck above.

Assuming both raised floor heights and ceiling heights are within our tolerances, PTS often utilizes the tops of the RLUs as a primary cable pathway for most row and/or pod branch circuits. We reserve the raised floor for PDU power cabling and branch circuits for non-cabinet free-standing equipment not well served by overhead distribution.

Determine the Cabinet Power Distribution Strategy

With the proliferation of dual power supply IT infrastructure, the obvious choice is to provide dual branch circuits to the equipment. Ideally, the branch circuits come from different PDUs, UPSs, generator/ATSs, and utility sources for true 2N (systems + system) or 2(N+1) redundancy.

In the case of single cord loads, the traditional methodology is to provision static switched PDUs to provide a single output from two (2) separate input sources. Again, the goal is to achieve 2N or 2(N+1) redundancy. A more recent approach is to provide rack-mounted automatic transfer switches at each RLU that require dual power distribution for single cord loads.

Power Distribution Impact on Cooling

The choice of how to manage the power cables within the RLU begins with deciding where to mount the power strips. If they are mounted vertically, they should be done so in the rear and to the side. If horizontal power strips are utilized, they can be mounted at any U location that is convenient per the equipment layout. In either case, the key is to dress the cables such that they do not impede the exit air flow of the IT equipment in the RLUs.

As previously stated, effective cooling has two (2) criteria, the delivered air temperature and the volume of air. As such, it doesn’t matter how cold the air is if the server is blocked from supply and return air flow.

Determine the Room & Cabinet Data Cabling Distribution Impact

Typically, there are three (3) choices in delivering network connectivity to an RLU. They are;

1. Home run every data port from a network core switch
2. Provide matching port-density patch panels at both the RLU and the core switch with pre-cabled cross-connections between them, such that server connections can be made with only patch cables at both ends

3. Provide an edge switch at every rack, row, or pod depending on bandwidth requirements. This approach is referred to as zone switching.

Alternatively, a hybrid of all or some of the approaches can be utilized. For instance, although a zone switch might be used for a pod of RLUs, a cross-connect field may be provided at the edge switch such that final port connections can be made with patch cords alone.

Data Cable Distribution Impact on Cooling

Once again, the impact on cooling performance is directly attributable to the effectiveness of the cable management strategy. The more densely the servers are packed, so go the cables. The key is to utilize RLUs with ample cable management space to avoid blocking air flow.

It should be noted, while server cable management arms are very functional, those that severely impede air flow should not be used, for the reasons previously stated.

Establish a Cooling Zone Strategy

Recall that effective computer room cooling is as much about removing heat as it is about adding cold.

Generally speaking, the three (3) equipment cooling methods along with their typical cooling potential can be determined from the following table;

Room Cooling	~2 kW per RLU
Row Cooling	~8 kW per RLU
Cabinet Cooling	~20 kW per RLU

Table 2

Room cooling is the process whereby air conditioners are positioned such that they are primarily cold air delivery sources, as opposed to heat removal sources. An up-flow CRAC unit delivering cold air to the room such that it is mixing with the heat being generated by the load is a typical data center example.

Row cooling is the ever popular raised floor, hot aisle / cold aisle cooling configuration. In this approach, cold air is delivered via perforated tiles in the cold aisle. The heat builds up in the hot aisle and uses the ceiling height as a pathway back to the CRAC units. Additionally, row cooling can be accomplished utilizing a new CRAC type commonly referred to as an 'in-row' system. This approach is discussed more completely in the next section.

Cabinet cooling is when the air conditioning is delivered in a self contained RLU. This is similar to a refrigerator with servers in it. The caveat to cabinet cooling is that it requires connection to a remote heat removal source such as chilled water. Otherwise, the cabinet will contribute to higher room temperatures. Also, it means you need an air conditioner in every cabinet, which is a very expensive proposition. Finally, if a cabinet cooling system fails, the doors typically are automatically controlled to open. Therefore, cabinet cooling approaches require a portion of the cooling load, typically twenty-five percent (25%), be provided in a room or row-based approach to serve as a backup cooling source.

High-Density Cooling

One common theme in data centers today is the idea of ever increasing load densities. Cooling RLU densities greater than 4 kW per RLU, but less than 8 kW per RLU can best be accomplished with in-row air conditioner solutions. Two (2) excellent products on the market are APC's RC series and Liebert's XD series. Although their approaches are completely different, both accomplish high-density cooling by doing the same thing. Both, closely couple heat recovery from the hot aisle and deliver large volumes of cold air directly to the cold aisle. APC's RC unit is a 12" chilled water air conditioner that gets sandwiched between RLUs. APC's RP unit is a 24" air-cooled air conditioner that gets sandwiched between RLUs. Liebert's XD series is a similarly functioning system available in both chilled water and DX versions, but can also be mounted directly over the RLUs.

Zone Cooling

An excellent strategy is to utilize combinations of the above methods in a hybrid approach. The result is a computer room that provides cooling in zones of capability. Some can be low-density RLU zones that require only raised floor row cooling, while others can be in-row cooling zones to provide medium-density cooling capacity. Still others can be cabinet cooling solutions for ultra high-density server environments.

A few other techniques that can be utilized in a zone cooling strategy are:

1. Spreading the load
2. Borrowed cooling or air scavenging

In essence, they are the same thing. The intent is to place higher density loads directly adjacent to lower density loads, thereby stealing, borrowing, or scavenging the excess cold air.

In any case, provisioning cooling zones is most easily accomplished by matching the cooling techniques with the load requirements as a result of analyzing the CFD model.

Determine the Cooling Methodology

Upon determining what cooling zone will be required, the decision of what types of air conditioners will be needed, must be made. There are four (4) air conditioner types;

1. air cooled
2. glycol cooled
3. condenser water cooled
4. chilled water

Direct Expansion (DX) Systems

The first three types of air conditioner types are called direct expansion (DX) systems. This means they condition air as a direct result of compressing and expanding gas. When pressure is applied to a gas its temperature increases and vice versa.

An abbreviated description of the refrigerant cycle of air conditioning is as follows;

1. Refrigerant is warmed by the heat from the equipment (cold air blown into the room is a byproduct of this stage). It is then made hot by compressing it
2. The high-pressure, hot refrigerant is pushed to an outdoor heat exchanger (condenser, dry-cooler, fluid-cooler, or cooling tower). The refrigerant is cooled by releasing its heat into the atmosphere
3. The now warm refrigerant is then returned to the CRAC where it is depressurized, or expanded, and once again made cold. And the whole process is repeated.

Chilled Water Systems

The lone non-DX system is the chilled water air conditioner. A chilled water air conditioner does not utilize a compressor as you cannot compress water. In this system, chilled water (about 45 degrees F) is circulated in piping and coils and is used to absorb the heat from the equipment. Interestingly, water is the best medium for the absorption of heat. In fact, it can absorb considerably more heat than an equivalent measure of refrigerant. However, the stigma of water in the data center caused this method to lose favor upon the disappearance of large mainframes. But, chilled water air conditioning is once again on the rise with the stratospheric heat densities being realized in today's data center. In June 2006, Hewlett-Packard announced it was building an entire data center utilizing its new water cooled server technology.

Heat Rejection

As we have discussed, heat rejection is an important part of the air conditioning process. When the heat rejection equipment is located inside it is called a self-contained system. For instance, a portable air conditioner is a self-contained air conditioner. The supply air (cold air) is released into the room and the return air (hot air) is usually exhausted into the plenum above the drop suspended ceiling. In this case, the heat isn't really so much rejected as it is displaced. Conversely, when the heat rejection equipment is located outside (or inside, but vented outside) the system is described as a split system. It is referred to this way because the evaporator section (the cold supply air) is separate for the heat rejection system (the hot return air).

Many times the choice of what type of air conditioner to use is predicated by the facility services readily available. For example, if a computer room is located on the 10th floor of a 38-floor building serviced by a condenser water loop, the obvious choice is a water cooled system, assuming the building will provide access to it and it operates on a 7x24 basis. In any event, proper sizing, selection, and installation is best left to qualified professionals.

Cooling Redundancy

It is important to consider redundancy factors when selecting an appropriate cooling solution. Redundancy means providing additional capacity such that if a portion of the system fails the remaining working components can deliver the designed capacity. Additionally, it is important to provide redundancy for all of the air conditioner sub-systems including CRAC units, heat rejection equipment, piping, pumps, etc. Sometimes, it is not enough just to provide an additional CRAC unit. Consideration must be given to placement as well. Not every CRAC unit may have the ability to provide every piece of equipment in the computer room. This is another reason to define the space in terms of cooling zones.

Typically, PTS will describe a zone in terms of its redundant capacity. For instance, if there is a raised floor air requiring (3) 20-ton CRAC units for capacity, but the space is longer than it is wide, it may require two (2) units for redundancy. This can be confirmed by CFD analysis. We describe this scenario as a 3+2 or 2-for-5 zone.

Precision Air Conditioners Versus Comfort Cool Air Conditioners

A precision air conditioner is the type typically found in a computer room. A non-precision, or comfort cool air conditioner is typically used in residential and commercial office construction. The difference between the two lies in their ability to provide:

1. Higher sensible heat ratios (SHR)
2. Tighter temperature control (typically 1 ºF)
3. Humidity control

Sensible Heat Ratio (SHR)

Simply stated, sensible heat is the heat given off by electronic equipment. Electrical energy is dissipated not consumed. Every watt delivered to a piece of IT equipment is dissipated as heat into the local environment. Sensible heat can be thought of as 'dry' heat because it has virtually no moisture content or stated another way, it is the air temperature. Another form of heat is latent heat. Latent heat can be thought of as 'moist' heat as it is the measure of heat contained in water molecules. It is the heat given off by human beings in the form of perspiration.

Both computer room grade (precision) air conditioner equipment and comfort cool air conditioner equipment remove both sensible and latent heat. The ratio to which they do so is referred to as the sensible heat ratio. A precision air conditioner can provide as much as 95% sensible heat cooling as opposed to a comfort cool air conditioner which can sometime deliver as little as 30% sensible heat cooling. Therefore, comfort cool air conditioners are a poor choice for IT environments.

Humidity

Comfort cool air conditioners cool a space by stripping moisture out of the air (the heat is contained in the water droplets). This is bad for computer rooms. Ideally, data centers require being maintained at 45% relative humidity (R.H.). This prevents static build up (too dry) and corrosion (too wet).

Determine the Cooling Delivery Methodology

Different architectural attributes affect cooling performance in different ways. For instance;

1. Computer rooms located on outside walls of the building, or worse with exterior windows are in bad locations due to the difficulty the air conditioners will have in adapting to the wildly varying environmental conditions presented by outside temperature, condensation, sunlight, etc.
2. The height of the raised floor directly affects the cooling performance of the CRAC units. PTS's CFD analysis has shown that 18" is the ideal height for most raised floor cooling applications. If the plenum is too low, the air moves too quickly, too close to the CRAC to be effective. It may even suck air into the raised floor (wick affect) as it rushes by. If the raised floor is too high, the air moves too slowly and the CRAC equipment has to work too hard to build up static pressure.
3. The height of the suspended ceiling also directly affects the cooling performance of the CRAC equipment. If the ceiling is too low, the heat cannot work its way back to the CRAC equipment to be rejected by the outdoor equipment. As a rule-of-thumb, don't start with anything less than a 9' ceiling. Increase the ceiling height one foot (1') for every twenty (20) watts per square foot of white floor space and not less than eighteen inches (18") above the tallest piece of equipment. It should be noted, in-row air conditioners are an excellent choice for low ceiling environments.
4. Many often question the necessity of a raised floor. PTS almost always recommends using one when practical. We realize it may not be the appropriate technique for the highest density cooling zones, but the raised floor cooling approach remains the most predictable and adaptable way to distribute cold air. Additionally, it provides a very convenient way to provide effective fluid piping pathway, fluid draining, power and data cable pathway, and cable management.
5. Bypass airflow is the most prevalent cooling problem in raised floor computer rooms today. Bypass airflow is the air that escapes from the raised floor in areas not serving the load directly. It is wasted valuable cold air. It can be found at every cable penetration through the raised floor that is not properly sealed with a brush-grommet. Any air that does not flow exactly where it was intended is bypass airflow.
6. Any cold air that serves to cool the hot air, in the hot aisle is wasted cold air. The concept of 'hot-spots' within a computer room is widely misunderstood. There are supposed to be hot spots in computer rooms, they are called hot aisles. As long as the heat has an effective path back to the CRAC and the inlet air to the servers is appropriate, let them be.
7. Hot spots are better described as areas where the inlet temperature is higher than it should be. Often this caused by re-circulated hot air. If a server is delivered an insufficient volume of cold air, it is considered being 'starved' and will compensate by drawing whatever air is available. Most commonly this air is hot air from the hot aisle. When hot air from the hot aisle mixes with the cold air, the net affect is insufficient cooling. While truly fixing this condition involves providing an adequate cold air volume, one good practice is to provide blanking plates in every U that is not occupied by a piece of equipment. Maintaining separation between the hot and cold aisles is a crucial aspect of good cooling design.

8. Dedicated returns are a common technique used to improve the cooling performance of a CRAC system. The idea is to use the suspended ceiling plenum (the space between the ceiling tiles and the deck above) for hot air return. The purpose is to get the hot aisle hot air back to the CRAC units as quickly as possible. This is accomplished one of two (2) ways; 1. Provide duct work from the top return of the CRAC units directly into the suspended ceiling plenum, while at the same time, placing grated ceiling tiles directly over the hot aisles (this can be thought of as a reverse raised floor), or 2. Provide assisted duct work (fan controlled) from the grated ceiling tiles directly over the hot aisles to grates directly above the CRAC unit's hot air return. PTS has used this technique quite effectively to alleviate hot spot issues without having to add additional cooling capacity.
9. Caution should be exercised before embarking on the 1st approach. Once conditioned air is drawn into the ceiling cavity, the space is described as a plenum. This means that all of the structures in the plenum need to be plenum rated. This includes every light fixture, power cable, data cable, etc. Not adhering to the safety restrictions is an unsafe practice.
10. A more recent development in the battle against high-density demand is APC's 'hot aisle containment' system. Hot aisle containment involves building a box out of the hot aisle such that it traps the hot air from leaving the hot aisle. The air conditioners are placed directly in-row and capture the hot air and reject it appropriately. In this approach the hot aisle temperature can easily exceed 100 °F which is never recommended. At these temperatures, the chilled water air conditioner is highly effective at delivering massive amounts of cooling. However, this technique lacks confirmation of the long-term viability of this approach.

Determine the Floor Plan

The 'hot aisle / cold aisle' approach is the accepted layout standard for RLUs for good reason. It works. It was developed by, Dr. Robert Sullivan, while working for IBM and it should be adapted for both new and retrofit projects.

Another big layout decision is where to place the CRAC units for a raised floor, row cooling arrangement. Traditionally, the CRAC units were placed at the perimeter of all four walls of the room. It was eventually realized the actual cooling performance was not matching that of the predicted performance. The reason has to do with fluid dynamics. Air streams perpendicular to one another collide and cause turbulence which negatively affects static pressure, which negatively affects cooling performance.

The next evolution placed the CRAC units on the opposing walls. This works and is still a common practice. However, it produces varying static pressure at different places on the floor. In the middle of the floor, where the two air streams collide head-on, is an area of very high static pressure. This drops off as you get closer to the CRAC equipment. Some data center professionals prefer to design for an even static pressure at all places of the floor to provide the same predictable amount of cooling to all equipment.

At PTS however, we are comfortable with a varying air flow raised floor. The reality is a computer room rarely has evenly distributed and similar loads. CFD analysis provides enough data to provide a predictable yet varying amount of cold air at every point in the floor. Furthermore, the CFD analysis will make obvious what zones can be used for low density, medium density, etc.

Common obstacles to effective cooling layout design are space constraints. Unfortunately, the design is often forced to accommodate a limited availability of space. In this case, experience is the best ally and the best approach is to consult a design expert.

Establish Cooling Performance Monitoring

It is vital to develop and deploy an environmental monitoring system capable of monitoring each room, row, and cabinet cooling zone. A given is that once effective cooling performance is established for a particular load profile, it will change rapidly. It is important to compile trending data for all environmental parameters for the site such that moves, adds, and changes can be executed quickly. At a minimum, the sensors that should be deployed are as follows:

1. At least one (1) equipment inlet air temperature sensor per RLU located at or near the top of the RLU where it is warmest. This should be done for every RLU in each distinct cooling zone and/or row of equipment.
2. One (1) to two (2) equipment outlet air temperature sensors per row located at or near the top of the RLU where it is warmest. This should be done for each distinct cooling zone and/or row of equipment.
3. One (1) or two (2) general room humidity sensors. Often it is sufficient to couple these with the RLU-based temperature sensors. Humidity can be room generalized because moisture disperses evenly throughout a given space regardless of air flow dynamics.

In the past, when facility departments had ultimate domain over all air conditioning, non-IP-based protocols were used to establish site monitoring. These systems are called building management systems. While they are expensive to deploy, the data they compile is extensive and can cover all facility support infrastructure. However, the translation of the data center portion of this information into TCP/IP format is cumbersome at best. Recent trends are seeing facility monitoring systems embracing the network.

PTS will typically deploy an IP-based, GUI monitoring system unique from that of the building systems. In lieu of monitoring the actual CRAC and heat exchange equipment, it is best to monitor the actual equipment inlet temperature for each cooling zone. This allows alerting to specific conditions that require attention regardless of the source of the problem.

Approaches Not to Take

There are a number of commonly deployed cooling techniques that should not be implemented. They are;

1. Reducing the CRAC supply air temperature to compensate for hot spots
2. Using cabinet and/or enclosures with either roof-mounted fans and/or under-cabinet floor cut-outs, without internal baffles
3. Isolating high-density RLUs

Reducing CRAC Temperatures

Simply making the air colder will not solve a cooling problem. The root of the problem is either a lack of cold air volume to the equipment inlet or it is lack of sufficient hot return air removal from the outlet of the equipment. All things equal, any piece of equipment with internal fans will cool it self. Typically, equipment manufactures do not even specify an inlet temperature. They usually provide only a percentage of clear space the front and rear of the equipment must be maintained to ensure adequate convection.

Roof-mounted cabinet fans

CFD analysis conclusively proves that roof-mounted fans and under-cabinet air cut-outs will not sufficiently cool a cabinet unless air baffles are utilized to isolate the cold air and hot air sections. Without baffles, roof-mounted fan will draw not only the desired hot air in the rear, but also a volume of cold air from the front prior to being drawn in by the IT load. This serves only to cool the volume of hot air which we have previously established as a bad strategy. Similarly, providing a cut-out in the access floor directly beneath the cabinet will provide cold air to the inlet of the IT loads, however, it will also leak air into the hot aisle. Again, this only serves to cool the hot air.

Isolating high-density equipment

While isolating high-density equipment isn't always a bad idea, special considerations must be made. Isolating the hot air is in fact, a good idea. However, the problem is in achieving a sufficient volume of cold air from the raised floor. Even then, assuming enough perforated floor tiles are dedicated to provide a sufficient air volume, too much of the hot air re-circulates from the back of the equipment to the front air inlet and combines with the cold air.

Conclusion

In conclusion, the first step in designing an effective cooling system is to establish a set of goals and criteria to serve as a basis of the design. The design criteria should define an adaptable, highly-available, maintainable, manageable, and cost-effective solution.

The following is a summary of the previously discussed 'best practices' in the chronology of how the design and development should proceed:

1. Determine the critical heat load
2. Establish the critical loads in terms of watts-per-RLU
3. Determine the CFM requirement per RLU
4. Identify the cooling performance impact from the room's power distribution strategy
5. Identify the cooling performance impact from the RLU power distribution strategy
6. Identify the cooling performance impact from the room & RLU data cabling strategy
7. Divide the room into cooling zones by RLU
8. Determine the most appropriate air conditioner type(s)
9. Determine the most effective cooling delivery methodology(s)
10. Establish a floor plan
11. Deploy a comprehensive IP-based monitoring system

About the Author:

Pete Sacco is the founder and President of PTS Data Center Solutions, Inc. (PTS). Pete has a BSEE and has been involved in the data center / computer room design industry for over 10 years.

About PTS Data Center Solutions

PTS Data Center Solutions is a data center consulting, design, and engineering services firm, as well as turnkey solutions provider.

We offer a broad range of project expertise. We specialize in planning, designing, constructing, monitoring, and maintaining computer rooms that integrate ‘best-of-breed’, critical infrastructure technologies. The result is an always available, scalable, redundant, fault-tolerant, manageable, and maintainable data center environment.

In today's hyper-competitive markets, where you measure network downtime in lost profits, PTS protects against some of the leading causes of downtime, hardware damage, data loss, and decreased productivity. PTS sets the standard for ‘always available’ solutions for data centers, computer rooms, server rooms, network closets, telecommunications rooms, network operations centers, and other mission critical facilities.

From our corporate headquarters in Franklin Lakes, New Jersey, and our office in Orange County, California, PTS works to fulfill our mission of creating satisfied customers by emphasizing pre-design & planning services to provide the optimal solution to meet our clients’ needs. This results in an early & accurate alignment between scope, schedule, and budget.

Experts for Your Always Available Data Center

568 Commerce Street
Franklin Lakes, NJ 07442
(201) 337-3833
(201) 337-4722 fax
www.PTSdcs.com