SITE INFRASTRUCTURE

WHITE PAPER

Reducing Bypass Airflow Is Essential for Eliminating Computer Room Hot Spots

By Robert F. Sullivan, Ph.D, with Lars Strong P.E. and Kenneth G. Brill.

Why is achieving reliable, stable, and predictable cooling in computer rooms so challenging?

Based on more than 15,000 individual measurements in 19 computer rooms ranging in size from 2,500 square feet (2,500 ft^2 or 230 m^2) to 26,000 ft^2 (2,400 m^2) and totaling 204,400 ft^2 (19,000 m^2) of raised floor, this study found that 10% of racks had air intake conditions outside the environmental parameters recommended by hardware manufacturers for maximum reliability and performance.

As hardware heat densities continue to increase, hot racks or cabinets will increasingly be a serious threat to not only information technology (IT) reliability and performance, but also to the willingness of hardware manufacturers' to honor fixed price maintenance contracts.

This pioneering study determined that high-heat densities and/or inadequate cooling capacity were not the underlying cause of hot spots. The highest percentage of hot spots was found in computer rooms with very light loads. Additionally, between 3.2 and 14.7 times more cooling capacity was running in those rooms than was required by the actual heat load. This white paper outlines counter-intuitive findings and relatively inexpensive solutions to recover wasted cooling capacity.

In this white paper

- Sixty percent of the available supply of cold air in the computer rooms studied is short cycling back to the cooling units. Called "bypass airflow," this means that only 40% of the cold air supply is directly cooling computer equipment. The remaining 60% of cold air mixed with the exhaust air is exiting from the heat load. This un-engineered mixing of ambient air provides indirect and uncontrolled cooling, especially for the equipment at the top of racks.
- Bypass airflow occurs through unsealed cable cutout openings and mis-located perforated tiles placed in the hot aisle. Based on a 10,000 ft^2 (930 m^2) computer room, bypass airflow averaged 80,000 cubic feet per minute (80 kcfm or 2,300 m^3/min) of which 31 kcfm (900 m^3/min) is due to misplaced perforated tiles and 49 kcfm (1,400 m^3/min) is due to unsealed cable cutout openings. This 80 kcfm (2,300 m^3/min) of under utilized cold air is the airflow equivalent of nine typical computer room cooling units.
- Recovery of wasted bypass airflow simply requires relocating misplaced perforated tiles and sealing unmanaged cable cutouts.
- In an operating data center, remediation should be undertaken taken only after completing a baseline study of how cooling is actually occurring. Closing too many openings in the wrong order or too quickly, may dramatically upset existing ambient cooling conditions and result in unintentionally overheated computer equipment.
- In one case study, bypass airflow was reduced from 43% to less than 10%. The hottest cabinet intake temperature dropped from 86°F (30°C) to 70°F (21°C). Even more importantly, the associated relative humidity increased from a dangerously low 20% up to 40%, which is the minimum value recommended by computer manufacturers for reliable computer operation.
- In a second case study, serious hot spots were eliminated when 11 out of 24 cooling units were turned OFF. Total data center energy consumption went down by 25%! The fact these savings were two times greater

In this white paper (cont.)

than had been forecasted indicate other significant performance factors were also in play.

- Remediation savings for a typical 10,000 ft^2 (900 m^2) computer room ranged from $85,000 to $110,000 per year (depending upon kWH rate) while cooling stability and reliability also improve.
- This research found that Computational Fluid Dynamics (CFD) modeling should used only after bypass airflow has been reduced below 10% and then only after values assumed in the model for cooling unit airflow and sensible cooling have been field measured.
- Before accepting the validity of CFD model results, actual airflow and temperature measurements must be rigorously and systematically compared to the model's predicted values. Gaps must be narrowed before making engineering or management decisions based upon the model.
- When used in conjunction with The Uptime Institute's® (the *Institute*) guidelines for Cold Aisle/Hot Aisles, the principles outlined in this white paper allow predictable air distribution using existing raised floor cooling technologies (25% open perforated tiles and downflow mounted cooling units) for heat loads of 4 kW per rack or cabinet over large areas. With careful engineering and management discipline to fully implement the *Institute's* other best practices, localized spot loads may significantly exceed the overall average.

Introduction

Engineers from the *Institute* recently completed a comprehensive survey of actual cooling conditions in 19 computer rooms comprising 204,400 ft^2 (19,000 m^2) of raised floor. More than 15,000 individual pieces of data were collected. Triton Technology Systems, Inc.® sponsored this research.

What follows in this white paper are valuable excerpts from this original and important research into consequences of current computer room planning and cooling practices.

This project required a total of 2½ years of original work with many of the measurements being made in 2002 through 2004. With the recent increase in computer manufacturer shipments of high-heat density products, the *Institute* strongly suspects today's percentage of racks with hot spots would be significantly greater than what is being reported in this paper. The authors believe the findings in this white paper are broadly representative of computer rooms in the U.S.

Prior to the original research in this paper, the actual operation of the air delivery portion of raised-floor computer room cooling systems was poorly understood. If a computer room was too hot, the tendency was to add more capacity. As shown in this research, overcapacity can actually make hot rooms even hotter, not colder, while incurring significant incapacity and cost penalties plus adding unnecessary reliability risks. We also conducted investigations into temperature "plumes," both under and above the raised floor. Plumes become an extremely important consideration if the gross remediation steps outlined in this paper don't fully resolve hot spot problems.

Gathering the data underlying this white paper required development of field measurement techniques to quantify what was really happening. Finding ways to consistently and accurately quantify the volume of air flow in a non-laboratory setting was difficult and delays occurred as actual findings divergent from original hypotheses were analyzed. Much of what was found was counter-intuitive. The authors found numerous instances of significant differences between what was actually measured and published technical data. The data set collected is much more comprehensive, and issues discovered are more detailed than what can be discussed here.

Hot Spots Come in Two Varieties

Hot spots occur in two varieties, zone and vertical. Zone hot spots occur when the temperature at all air intake levels of a rack or cabinet are too hot. Zone hot spots typically exist over large areas of the raised floor. Vertical hot spots are more discrete and may be unique to just a single rack or cabinet. Identifying a vertical hot spot requires measuring the ambient temperature vertically up the intake face of a rack or cabinet. An abrupt temperature change (5°F to 15°F or 3°C to 8°C) will occur over a short vertical distance (6 inches or 15 centimeters) revealing a vertical hot spot. The lower the vertical height at which this transition occurs, the more serious the hot spot problem.

Vertical hot spots occur because the internal fans within the computer equipment at the bottom of a rack or

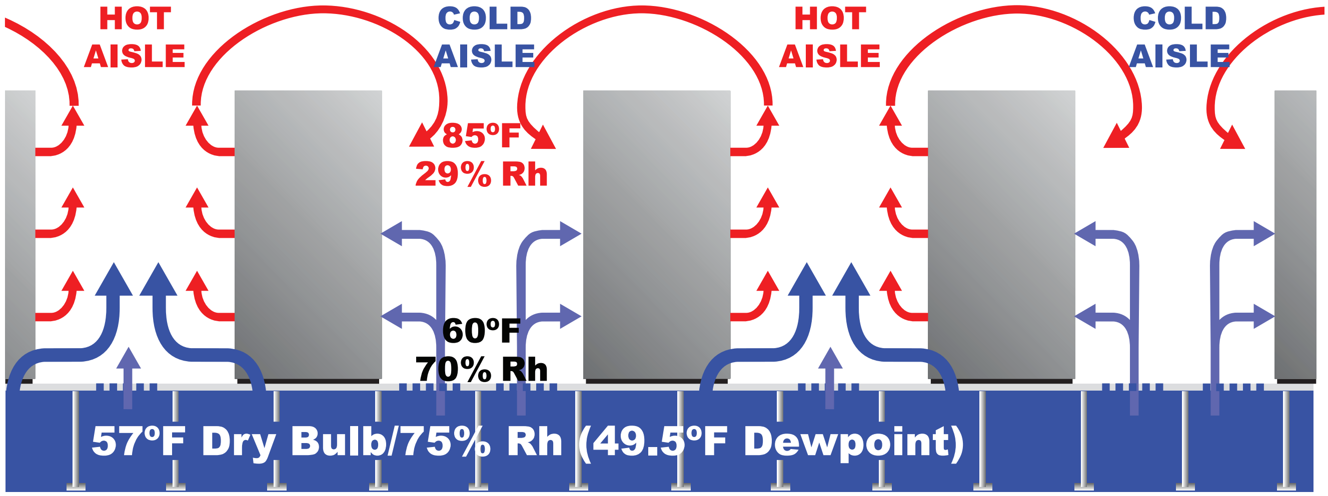

Figure 1. Recirculation of hot exhaust air across the top of racks due to an inadequate supply of cold air from perforated tiles will result in unacceptably high intake air temperatures for the equipment housed in the top of racks. (The majority of the hot racks reported in this paper were hot only at the top of the rack.)

cabinet have consumed the available supply of cold air coming from nearby perforated tiles. With no cool air remaining, equipment above the temperature transition is left to pull air from the hot exhausts of adjacent computer equipment or from the ambient conditions in the room.

Hot Spots Are a Serious Threat to Maximum Information Availability and Hardware Reliability

Virtually all high-performance computer, communication, and storage products now incorporate internal thermal sensors that automatically will slow or shut down processing when temperatures exceed predetermined thresholds. Achieving and maintaining high availability requires that these sensors never be triggered by customer supplied and controlled environments.

Computer and communication hardware manufacturers have collectively published through the American Society of Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE) and the TC 9.9 Technical Committee, their recommended environmental conditions (air intake temperature and relative humidity) on the performance of their equipment. Two performance classifications are defined by the manufacturers, the first being whether the equipment will operate (i.e., it will run, but perhaps not with maximum reliability) and second being the conditions required for maximum performance and reliability. These manufacturer recommendations are shown in Figure 2 along with the *Institute's* recommended cooling unit return air control set-point.

A number of very important points can be drawn from Figure 2:

- Environmental conditions are determined at the air intake of the equipment. The discharge temperature of the hardware exhaust air, or the air measured 48" (1.2 m) above the floor, or back at the cooling unit is of little concern. What counts, in terms of reliability, performance and warranty, are the conditions at the equipment air intake.
- While the equipment may operate at air intake temperatures of up to 90°F (32.2°C) or at a relative humidity within 35% to 80% (depending upon dry bulb), it may not run reliably or at specified performance standards.
- For maximum performance and reliability, computer manufacturers recommend a maximum temperature of less than 77°F (25°C) with a rate of change not exceeding 9°F (5°C) per hour.
- For maximum reliability and performance, relative humidity (Rh) must exceed 40% with a rate of change not to exceed 5% per hour. The threat of spontaneous electrostatic discharge begins to occur when Rh is 30% or less.
- Users can't control rack intake air temperatures which may be 50 feet (15m) away from cooling units. What they can control is the return air temperature and relative humidity back at cooling units.
- The divergence between computer manufacturer air intake requirements and what users can control and deliver is a major industry problem.
- The *Institute's* recommended return air control point at the cooling unit is 72°F/45% Rh (22.2°C/45%

Figure 2:
Computer Hardware Reliability Environmental Reliability Limits

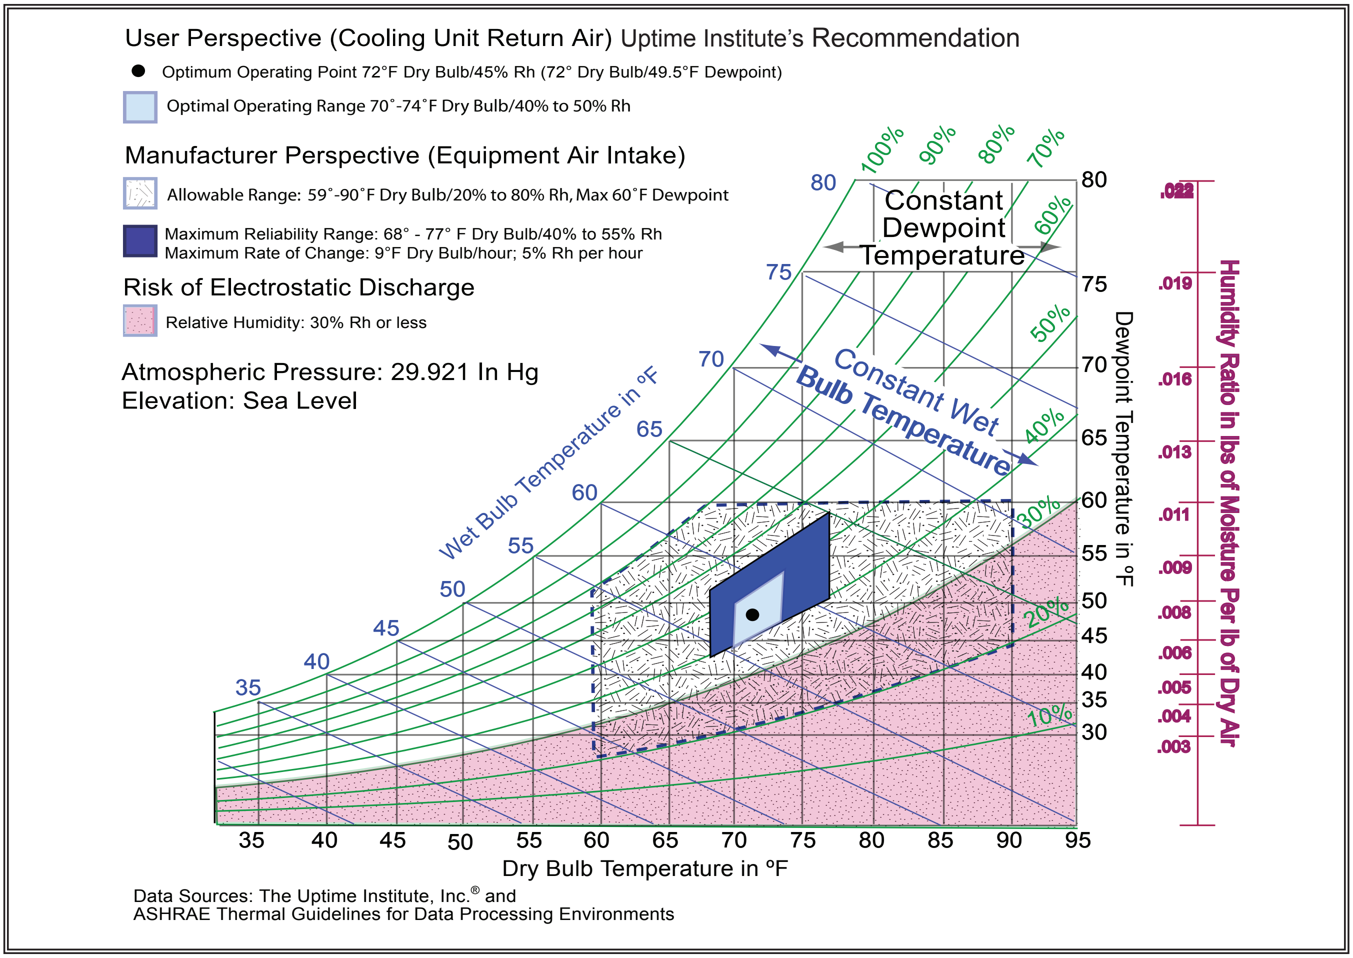

Rh). This allows for optimal cooling unit efficiency with some tolerance for local hot spots.

High Soft Error Rates, Erratic or Unrepeatable Information, and Outright Hardware Failures Can Result from Exceeding Recommended Environmental Recommendations

Excessive expansion and contraction of the component materials inside computer and communication equipment results from exceeding manufacturer temperature range and rate of change recommendations. When the equipment in a computer room is "cooked," a considerably higher rate of premature equipment failure can be expected over the following days, weeks, and months, even though there were no failures during or immediately following the event. While failures may be instant, more likely they will be latent and may take days or weeks to appear. During the intervening period, apparently functional but damaged, equipment is creating reliability ghosts and loss of processing ability, which frantic technicians are tearing their hair out trying to isolate and correct.

The simplest example of thermally caused instability is when the electrical contacts of a printed circuit board which plugs-on to an interconnecting wiring backplane no longer make physical contact resulting in intermittent or outright failures. Experienced technicians know the first thing to do after a thermal excursion or thermal cycle[1] is to re-seat all cards in the card cage. This often solves erratic operations by restoring positive electrical connections.

Another example of thermally caused performance issues is a high soft error rate. The hardware keeps operating, but at dramatically reduced speeds. This can

[1] A thermal cycle results from an electrical power down (planned or unplanned) during which the equipment has a chance to completely cool down. This allows maximum contraction. When the device is powered back up, maximum expansion occurs.

happen on disc drives when expansion or contraction has shifted where the data is physically located on the media. This results in reduced read/write throughput as multiple attempts need to be made to access the right information.

A final example is when thermal expansion has been so great that microscopic printed circuit traces that carries internal signals actually breaks. While outright failure may not occur at that exact moment, a ticking time bomb is created. The ultimate trigger for outright failure may not transpire for several months after the event or until the device is powered down for maintenance. Experienced technicians know that after a power down, an unusually high number of devices won't re-start. This is why experienced data center managers require a full complement of customer engineers and spare part from each manufacturer on site when a planned electrical infrastructure shutdown is planned.

At one Fortune 100 site, the most critical application experienced an availability failure despite millions of dollars invested in mirroring and extensive hardware redundancy. The application failure occurred within six weeks after a catastrophic cooling failure that could not be repaired quickly. A management decision was made to open the computer room doors to the outside environment and to continue operating the computer equipment. Ambient temperatures in some areas of the computer room exceeded 95 °F (35 °C) and relative humidity was uncontrolled. Information availability was successfully maintained until the cooling problem could be repaired. However, during the next six weeks, hardware failure rates exceeded normal field experience by more than four times. Information availability was maintained due to extensive redundancy, except for one hard failure when a second device failed before the first failure could be repaired. This caused the unscheduled system outage. This example illustrates that while there is often not an immediate connection between a thermal excursion and subsequent hardware fall-out, the circumstantial evidence for a direct cause and effect connection is extremely strong.

Excessive Hot Spots Are Already Present in Most Computer Rooms

Ten percent of the racks (10 of every 100 racks) in 13 computer rooms had ambient temperatures of 75°F (24°C)[2] or higher at the computer equipment air intake at the top of the rack.

Hot Spots Are Not Caused by Inadequate Cooling Capacity or High Heat Densities

For the 13 rooms studied in greatest detail, inadequate cooling capacity, high watts per square foot (W/ft^2), and the racks with intake temperatures of 75°F (24°C) or more could not be correlated. In both Computer Rooms 3 and 8 of Table 1, 15 times more cooling was running than was required by the actual heat load.[3] Despite having 15 times more cooling running than was required, one room had 20% hot racks/cabinets and the other had 7% hot racks/cabinets. Similarly, there was no relationship between hot racks and W/ft^2. The heat load in Computer Room 3 with 20% of the racks being too hot was only 4 W/ft^2 (43 W/m^2). Room 8, with 7% hot racks and had only 3 W/ft^2 (32 W/m^2) of heat load.

Excessive Bypass Airflow Is the Underlying Problem

The reason so much cooling overcapacity failed to cool such low density loads is that 59% of the cold air is bypassing the air intakes of the computer equipment. (59% is the weighted average of 13 rooms in Table 1 with Computer Room 5 being the best case at 34% and Computer Room 2 being the worst at 79%). What this means is that on average only 41% of the cold air is directly cooling computer equipment. With so little cold air going into air intakes, heat removal is actually occurring by the uncontrolled mixing of escaping cold bypass air with hot exhaust air.

"Bypass airflow" is defined as conditioned air that is short cycled or not getting directly into the computer equipment. This air escapes through cable cutouts, holes under cabinets, or misplaced perforated tiles. Bypass airflow occurs on the hot air discharge side of the computer equipment and pre-cools the hot exhaust air returning to the cooling units.

The larger 19-room study of 204,400 ft^2 (19,000 m^2) basically found the same bypass phenomena; namely, the weighted average bypass airflow was 60% (versus 59% for the smaller study). Of the wasted airflow, 39% was escaping through perforated tiles not in the cold aisle. The other 61% was escaping through unsealed cable cutout openings under racks and cabinets. For a 10,000 ft^2 computer room, the 60% bypass airflow amounted to a total of 80 kcfm (2,300 m^3/min) in cold air losses with 31 kcfm (900 m^3/min) escaping through misplaced perforated tiles and 49 kcfm (1,400 m^3/min) escaping through unsealed cable openings. Bypass

[2] A dry bulb warning threshold temperature of 75°F (23.9°C) was selected for this study because for a computer room controlled to a dew point of 49.5°F (9.7°C), Rh falls below 40% at a dry bulb temperature of 75.5°F (24.2°C). Relative humidity below 40% results in conditions susceptible to spontaneous electrostatic discharge.
[3] Required cooling was calculated by converting UPS power consumed to heat load and then adding 20% for cooling unit redundancy and 10% for bypass airflow.

Cooling Effectiveness in 13 Computer Rooms Comprising 170,000 Ft² (15,800 m²)							
Computer Room	Electrically Active Racks or Cabinets with Air Intake Temperatures ≥75°F (24°C)	Heat Load Density Across the Gross Computer Room W/ft²	W/m²	Excess Cooling Capacity Running	Cooling Unit Average Delta T °F	°C	Bypass Airflow
1	28%	8.7	94	5.2 times	5.3	2.9	63%
2	21	17	183	3.2	6.6	3.7	79
3	20	4	43	14.7	2.1	1.2	74
4	16	43	463	0.1	14.6	8.1	44
5	13	12	129	1.5	5.5	3.1	34
6	12	31	334	2.8	5.8	3.2	45
7	8	11	118	1.8	9.3	5.2	55
8	7	3	32	15.3	1.7	.9	40
9	4	8	86	5.2	4.0	2.2	53
10	3	31	334	1.9	6.5	3.6	75
11	1	8	86	4.0	4.9	2.7	71
12	1	13	140	1.0	6.0	3.2	67
13	1	33	355	0.3	11.5	6.4	61
Weighted Average	**10%**	**14**	**151**	**2.6 times**	**7.6**	**4.2**	**59%**

Table 1. For the computer rooms evaluated, vertical hot spots are more directly related to Bypass Airflow than excessive heat load density or inadequate cooling capacity

airflow can also escape through holes in computer room floors, walls, or ceilings, but this was not found to be significant in the rooms studied.

At the static pressures required to successfully cool 2 kW per cabinet, the average unsealed cable opening short cycles or wastes 1.9 kW of equivalent cooling. If these unmanaged openings were located on the equipment's air intake side and then the now unnecessary perforated tiles were removed, this waste would be of some cooling benefit. But with the cable openings on the exhaust discharge side of the equipment, each rack or cabinet position is wasting almost as much cooling as it consumes. From another perspective, each unsealed cable opening in this study short cycled the airflow equivalent of one half a perforated tile.

In contrast to the 2 kW per rack or 67 watts per gross computer room square foot (721 W/m^2) assumed in this paper's engineering calculations, currently-available fully configured blade or 1U high-end servers already consume an actual 8 kW or more per full cabinet/rack. New high-performance computer products to be released within the next 24 months will jump this to 16 kW or more per full rack or cabinet. Although water-cooling or another non-air cooling solution will eventually be required, some significant portion of the heat load will always be exhausted to the ambient air (in mainframe days, 85% of the heat was rejected to water and 15% to air – if the same ratio were to apply in the future and chilled water or other auxiliary cooling methods begin when heat loads reach 15kW per full cabinet, this means 2¼ kW per cabinet in heat will still be rejected to air).

By increasing airflow volume through perforated tiles, reducing or eliminating bypass airflow is already crucial, and will become more so in the future. Eventually 60% open grates will be required because perforated tiles don't have enough open area to allow sufficient airflow.

Figure 3. These misplaced perforated tiles and less obvious unmanaged cable openings are allowing valuable cold air to escape. Forty-five percent of the capacity of a 20-ton cooling unit is being wasted in this picture.

Bypass Airflow Causes Significant Inefficiencies

Reducing bypass airflow to 10% or less for the sites studied will reduce the cold aisle temperature by 4°F (2°C) to 8°F (4°C) while at the same time allowing for an increase in the cooling unit's return air temperature control setting. If the proper quantity of perforated tiles is installed in the proper locations, zone and vertical hotspots will disappear. This significant counter-intuitive cooling quality result also has the following significant savings in energy, maintenance, and other operating costs (capital costs are not considered in the following analysis)

- Reduced Fan Horsepower. From Table 1, 2.6 times more cooling was found to be running than the actual heat load required. This excess capacity was running purely to provide additional fan kCFM (m^3/min) capacity to compensate for lost bypass airflow. If bypass airflow was reduced to 10% or less by relocating perforated tiles and sealing currently unmanaged openings, 189 cooling units could be turned OFF at a savings of at least 5 horsepower or 4 kW each (cooling unit motors range from 5 to 7.5 with some being 15 horsepower depending upon unit type. Five horsepower was selected as the most conservative condition). At $0.06 per kilowatt hour (kWH), this reduction in horsepower consumption would amount to an annual savings of $1,960/cooling unit, $21,560 per 10,000 ft^2 (930 m^2) of computer room, or $370,440 for the 13 computer rooms. At $0.10/kWH, the annual savings would be $3,267/cooling unit, $35,940 per 10,000 ft^2 (930 m^2) of computer room, or $627,523 for 13 computer rooms.
- Reduced Maintenance. Maintenance would be eliminated on each cooling unit turned Off for a further savings of $300/month - $3,600 per year, $39,600 per 10,000 ft^2 (930 m^2), or a reduction of $680,400 per year for all rooms.
- Reduced Latent Cooling Penalty. In an effort to overcome local hot spots, many sites operate their cooling units with return air temperatures significantly below 72°F (22°C). Increasing return air temperature would eliminate or reduce significant amounts of de-humidification and re-humidification. This is both a cost savings and a significant IT availability risk reduction. Most water leaks under raised floors are related to humidification/de-humidification, and dust burning off reheat coils or humidifiers (which is the underlying root cause of many false fire alarms and gas discharge events).

Reducing the return air temperature of a typical 20-ton Liebert Computer Room Air Conditioner Model 267 W (DX type with water heat rejection) by just 2°F (1°C) from 72°F/45% Rh (22°C/45% Rh) to 70°F/48% Rh (21°C/45% Rh) reduces the sensible cooling capacity by 11% (from 229,000 British thermal units [Btu] per hour to 203 kBtu/hr. Moreover, the lost sensible tonnage is converted to latent cooling which then wrings moisture out of the air at the rate of 1.8 gallons per hour (6.8l/hr).[4]

If the return air temperature were raised by 2°F (1°C) back to 72°F (22°C), 1.8 gallons of moisture removal per hour (6.8 l/hr) would be eliminated (assuming the Liebert Model 267W cooling unit above). The energy savings is 4.8 kW per gallon (1.3 kW/liter) not removed. At $0.06/kWH this is $4,541/unit/year and at $0.10/kWH it is $7,570/unit/year. This does not include an additional savings from eliminating maintenance or replacement of humidifier components.

[4] In a 1 year period, 1.8 gallons per hour (6.8 liter per hour) removed plus an additional 1.8 gallons (6.8 l/hr) of moisture which must subsequently be added back into the computer room to maintain constant relative humidity amounts to 31,500 gallons (119,000 liters) for each cooling unit or enough to fill the 260 ft^2 (24 m^2) area of raised floor area served by the cooling unit at 67 W/ft^2 (721 W/m^2) to a height of 4.5 feet (1.4 m). No wonder water leaks in computer rooms are so frequent!

Assuming the return air temperature control point on a third of all cooling units is set to 70°F (21°C)[5] and unnecessarily removing 1.8 gallons per hour (6.8 l/hr), the reduction in the latent cooling penalty at $0.06/kWH would be $4,541 per cooling unit, $23,239 for a 10,000 ft² (930m²) room, and $592,304 for all computer rooms.

The total annual savings from reduced fan horsepower, reduced maintenance, and reduced latent cooling penalty (summarized in Table 2) by reducing bypass air flow from 60% to 10% is measured in millions of dollars! This optimization is truly unique because better cooling quality, greater reliability, and increased stability are significantly less costly than the results produced by current computer room cooling practices.

Remediation Case Study One (Success)

One of the computer rooms in the larger 19-room study consisted of 2,500 ft² (230 m²). Within this room, seven servers in the mainframe area had intake temperatures above 75°F (24°C) with one being as high as 86°F (30°C). In the Wintel server area, six racks had temperatures above 75°F (24°C) with the highest being 82°F (29°C). The return air temperatures at the cooling units ranged from 70°F to 72°F (21°C to 22°C) with the relative humidity at these temperatures being 35% to 40%. The cooling unit temperature controls were set at 66°F to 68°F (19°C to 20°C) in an attempt to reduce the hot spot temperatures. The relative humidity in the hot spot areas was 20% to 25%, well down into the electrostatic discharge risk area. Perforated tile airflow averaged 290 cfm and ranged from 198 cfm to 330 cfm (averaged 8.2 m³/min and ranged from 5.6 m³/min to 9.3 m³/min). Bypass airflow was measured at 43%.

Bypass airflow was reduced to less than 10% without shutting down computer operations by closing unmanaged cable openings with self-sealing Koldlok® products made by Triton Technology Systems. Inc®. Thirty-two cable openings under computer hardware were closed using standard KoldLok parts and filler plates. Openings larger than 4” by 8” (10 cm by 20 cm) and holes in perimeter walls were sealed with custom fabricated KoldLok parts that could be assembled in tight spaces with limited access height and where an entire floor tile was missing. Openings around six telecom racks, two patch panels, and two significant 3-inch (8 cm) gaps around cooling units were sealed using Koldlok extended assemblies. This project was completed over a period of four days without a computer equipment shutdown.

Upon completion of remediation, average airflow through the perforated tiles increased 81% to 526 cfm (15 m³/min). The highest hot spot temperatures dropped by 13°F, 14°F, and 16°F (respectively 7°C, 8°C, and 9°C) bringing all air intake temperatures well within the tolerance window recommended for maximum equipment reliability and performance. As a result, the worst relative humidity rose from 20% to 40%, which was within computer manufacturer recommendations. Cold aisle temperature reductions of 4°F to 6°F (2°C to 3°C) were common. Temperatures in the hot aisle

Table 2: Annual Savings From Bypass Airflow Reduction

Savings Source	Per Cooling Unit		Per 10,000 Ft² (930m²)		Per All 13 Rooms	
	$0.06/kWH	$0.10/kWH	$0.06/kWH	$0.10/kWH	$0.06/kWH	$0.10/kWH
Reduced Motor Hp	$1,960	$3,267	$21,560	$35,940	$370,440	$627,216
Reduced Maintenance	3,600	3,600	39,600	39,600	680,400	680,400
Reduced Latent Cooling Penalty	4.541	7,570	23,239	34,841	395,067	592,304
Total	$10,101	$14,437	$84,399	$110,381	$1,591,129	$1,899,920

Table 2. Summary of annual operating expense savings resulting from reducing bypass airflow allowing unnecessary cooling units to be turned OFF.

[5] This is a very conservative assumption as many cooling units are found set at 68°F (20°C) or even lower.

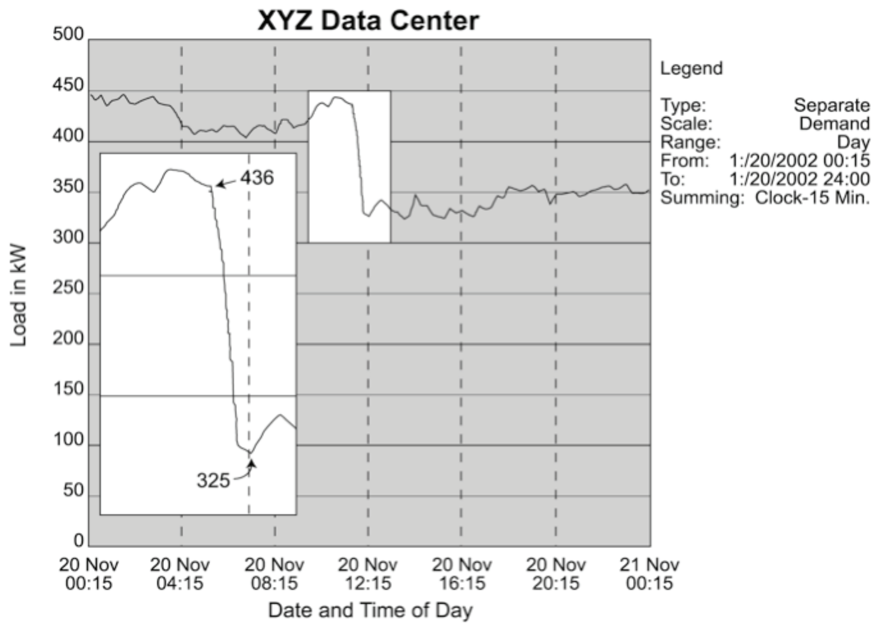

Figure 4. A worried utility manager faxed this electricity consumption graph to his data center customer noting an abrupt drop in energy consumption. The 111 kW per hour reduction far exceeded the kW saved from turning OFF eleven cooling units or approximately 55 fan horsepower. This indicates the reduction in de-humidification and re-humidification was a very significant factor in saving energy.

remained approximately the same. This surprising result is explained because as the uncontrolled mixing of bypass cold and hot exhaust air was eliminated, the temperature in the cold aisle dropped to approach the temperature of the supply air. Another benefit occurred when the level of audible noise in the room dropped significantly, but this was not quantified because the necessary instrumentation was not on hand. These cooling performance improvements occurred as bypass airflow dropped from 43% to less than 10%.

Although no cooling units were turned OFF, successful reductions in high air intake temperatures halted plans to install additional cooling units. The result was a $60,000 savings in previously planned capital expenditures.

Remediation Case Study Two (Success)

Computer Room 11 in Table 1, comprising 18,800 ft^2 (1,750 m^2), had very severe vertical hot spot problems despite having only 144 kW of heat load and 24 cooling units running (based purely on heat load, only 41 tons of cooling or (3) 20-ton cooling units were required to be running). After a thorough baseline study, 11 cooling units were turned off and the hot spots disappeared completely. Energy consumption at the electric utility meter went down by 25%! The local electric utility manager sent a worried fax asking what was happening. See Figure 4.

Energy consumption for the 11 cooling unit fans turned OFF (another 7 to 9 units could have been turned OFF if bypass airflow had been reduced from 71% to 10%) would have been about 41kW, yet total energy consumption went down by 111 kW per hour (or a annual saving of almost 1,000,000 kWH). The only explanation for the additional 70 kW reduction was energy consumed by dehumidification, re-humidification, and re-heat was eliminated as all other load and environmental conditions remained unchanged.

Remediation Case Study Three (Success)

This new computer room comprising 15,000 ft^2 (1,400 m^2) of raised floor was master planned with all future IT equipment locations identified in advance. The resulting IT hardware yield was 32 ft^2 (3 m^2) per rack or cabinet position. During commissioning, 500 to 700 CFM (14 m^3/min to 20 m^3/min) of airflow, 2.4 kW to 3.3 kW per perforated tile, or gross 80 W/ft^2 to 106 W/ft^2 (800 W/m^2 to 1,140 W/m^2) of cooling capacity was measured through any perforated tile on the raised floor. Bypass airflow was measured at 13%.

Remediation Case Study Four (Failure)

Several minutes after an arbitrary rearrangement of 30 perforated tiles, 250 servers automatically thermaled off due to internal safety controls within the hardware that are intended to prevent overheating. Internet service for a critical Application Service Provider was halted during prime time.

Remediation Case Study Five (Failure)

A mechanical system diagnostic assessment was performed in a 40,000 ft^2 (3,700 m^2) computer room with over 70 chilled water cooling units. This study found two cooling units had been piped incorrectly (supply and return were reversed) during original construction. Three units were found to have stuck chilled water valves. None of the five units annunciated an alarm indicating that their ability to cool had been compromised.

Using theoretical calculations, only 35 units (including redundancy) were required to keep the room cool. Upon receiving the study, the facility staff began turning cooling units OFF in an uncontrolled manner. Within five minutes, critical computer and DASD storage equipment had thermaled OFF (internal protective sensors had sensed high temperatures and had automatically shut down the hardware), interrupting mission critical computing. The IT people arrived in great consternation to find out why availability had been lost during prime time in a global system with over 160,000 active online users. All cooling units were turned back ON and the hot spot problems created by impulsive action immediately went away. After this disaster which involved angry senior executives of the company, no further attempts were made at cooling optimization.

This example of a failed remediation attempt illustrates the importance of making changes carefully and gradually in conjunction with extensive temperature monitoring and having a rehearsed back-out plan if environmental conditions don't improve as expected. This failure generated the impetuous for the *Institute's* current research. It became clear that the interactions were much more complex, dynamic, and immediate than anticipated and the science for predicting what would happen was too primitive. Using the research and conclusions reported in this paper, this failure could have been a success with the owner saving more than a million dollars while also improving computer room environmental reliability and stability.

Predictions from Airflow Modeling Software often Don't Match Actual Measurements

Computational Fluid Dynamics (CFD) modeling has become very popular and companies spend significant sums creating very elaborate maps of predicted computer room static pressure, airflow and temperatures. Unfortunately, these maps have failed by very significant margins to reflect actual field measurements made by the *Institute* in this and other related research.

In one room (on which the owner had spent hundreds of thousands of dollars for modeling), the actual measured results were that 50% of the racks in one zone had equipment air temperatures in excess of 75°F (24°C). The highest temperature was 96°F (36°C). In addition, the relative humidity was 23% which was well into the range for electrostatic discharge. (The site had been experiencing a higher rate of hard drive failures than would be expected under normal field conditions.) The model predicted the maximum temperatures anywhere in the area would be less than 72°F (22°C) or well within the acceptable range for maximum reliability and performance. How could this happen? This site had read the *Institute's* white papers and had made a determined effort to reduce bypass airflow using "off-the-shelf" KoldLok parts. However, they failed to fully appreciate the significance of needing to seal <u>all</u> cable openings, including those under equipment and openings where the entire floor tile is missing in order to get bypass airflow below 10%. As the easy openings are closed, static pressure rises which pushes more airflow through the remaining unclosed openings. This zone had 160 unsealed openings which amounted to 36 ft^2 (3.3m^2). The lost cooling through these remaining openings was equal to 55 perforated tiles or the approximate capacity of three cooling units.

In addition to excessive bypass airflow, only one of the five cooling units in the zone was producing rated sensible capacity. Two cooling units were producing less than 50% rated air flow indicating their blower belts were slipping. The reason for incapacity of the other two cooling units was not be determined. (On average, 10% of the cooling units examined in the course of this research had failed with no alarm or indication of failure.)

With actual bypass airflow significantly greater than assumed, two cooling units with slipping belts, and with two cooling units failing to specified deliver delta T, this CFM model produced results that were seriously wrong! Management was falsely comforted by fancy temperature maps and failed to validate results with measurements of actual results. Unfortunately, this real case study repeats itself much too often.

Validate Conclusions before Depending upon CFD Modeling

The benefits of computational fluid dynamic modeling can be substantial, but it should only be used after bypass airflow has already been reduced to 10% or less. (If all the openings are correctly measured and properly input into the model, which is costly, the recommended course of action is most likely to close the openings as this is much cheaper than any other alternative for improving cooling performance.)

The two most common and biggest modeling mistakes are using the cooling unit manufacturers specified airflow volume and failing to fully, if at all, include the actual bypass airflow openings. Sites often use the manufacturers specified cooling unit airflow volume because they lack the analytical tools and education necessary to make field measurements of the actual airflow being delivered. In 90% of sites studied by the *Institute*, at least one of the installed cooling units was delivering less than the rated airflow volume. In some cases the delivered volume was 40% of the rated capacity.

Differences between specified and actual cooling unit airflow volumes, otherwise failed cooling units, and actual bypass airflow losses have a substantial impact on the model's predicted static pressure and therefore the perforated tile airflow, and IT equipment intake air temperatures. Managers should be aware of the consequences of GIGO – Garbage In, Garbage Out!

Until each parameter in the CFD model and its interaction with other variables is fully understood by the modeler, users should not risk dependence upon model results without first independently validating predicted results with actual measurements. Even when accurate input values are used, some "tuning" is often needed to narrow the gap between actual and predicted results. Without this tuning (which currently is more art than science), facilities managers solely using "off the shelf" CFM models are likely to be surprised and disappointed by the gap between predicted results and actual conditions.

Computational fluid dynamics modeling is an extremely valuable tool, but the results are accurate only as long as the input values match actual site conditions (GIGO) and all significant variables, sometimes including specialized tuning, are included.

Justification for Bypass Airflow Remediation

Solving air distribution problems in the computer rooms studied is fast, relatively inexpensive, especially relative to installing additional cooling equipment. Construction often takes months and has the risk of unintended downtime. In fact, installing more cooling equipment is likely to make existing problems worse. The cost of bypass airflow improvements is likely to be self-funding through reductions in cooling expenses. However, and much more importantly, fixing air distribution problems will actually improve IT reliability and stability and reduce the frequency of IT chasing ghosts and gremlins. This is one of the few cases in data center facilities where reliability, availability, serviceability, speed of solution, and cost are all aligned.

Remediation Warning

Changing the quantity or location of perforated tiles, or closing cable and other openings without first shutting down computer equipment while airflow changes are being made is a high-risk proposition. Since halting computing is usually not an option, non-stop remediation is necessary, but involves high risk. Risk occurs because, on average, 60% of available cooling comes from the mixing of ambient air in the overall room and not from the perforated tile openings in the cold aisle. Changing perforated tile locations too quickly or closing too many openings in the wrong sequence can result in very rapid increases in IT equipment ambient air temperature increase.

Significant modifications to airflow or cooling systems should not be undertaken without enlisting expert advice and guidance. Detailed plans, both implementation and back out, should be developed before any work is initiated.

Unlike electrical systems, cooling systems often behave in counter-intuitive, extremely non-linear ways. Affecting change in one area can cause significant costs or affects in another area as a result of unexpected dynamic interactions. Failure to appreciate and respect this interrelated behavior can result in significant consequences. With air turnover rates of almost once a minute, it is likely that sites ignoring this warning could encounter severe hardware damage before knowing temperatures had gotten out of control (see Remediation Failure Case Study Four for things to be avoided).

Conclusions

Computer manufacturers have codified an environmental window of 68°F to 77°F (20°C to 25°C) and 40% to 55% Rh as maximizing hardware reliability and performance. They have also initiated a mechanism for possibly not honoring maintenance contracts in the future if customer conditions are outside that window. The gap between what is required at the air intake of the hardware and what IT and facilities can actually control and deliver is substantial. In particular, the problem of controlling relative humidity to assure a minimum of 40% Rh at all elevations is likely to be significantly more difficult than controlling dry bulb temperature to be less than 77°F.

In the computer rooms studied, the primary cooling problem was inadequate air distribution, not a lack of thermal cooling capacity or excessive heat load. Every room had a significant excess of sensible cooling, yet all cooling units had to remain running in order to compensate for bypass airflow losses which averaged 60%. Excessive bypass airflow through unmanaged cable openings resulted in low static pressure. Lack of sufficient static pressure resulted in zone hot spots where there wasn't enough cold air within the zone and in localized vertical hot spots where the supply of cold air was fully consumed by the equipment in the lower part of the rack or cabinet.

Substantial excesses of cooling unit fan capacity is currently allowing many poorly engineered or arranged computer rooms to function, although at lower reliability and added cooling cost than would otherwise could be achieved.

To successfully cool heat loads greater than 2 kW per rack deployed in groups of forty or more racks (the equivalent of 67 watts per gross computer room square foot or 6 watts per gross computer room square meter), experience and engineering calculations require that bypass airflow be reduced to 10% or less. Out of the 19 rooms evaluated, only three were even remotely close to achieving this requirement (i.e. the rooms with the lowest bypass airflow - 20%, 35%, and 38%). Cooling conditions that were formerly left to chance must now become carefully engineered and controlled.

Reducing bypass airflow requires placing perforated tiles only in the cold aisles, closing unmanaged openings, and other significant, but relatively inexpensive changes in computer room layout and operating practices. (Other alternatives to remediation require the installation of additional cooling capacity which itself is a high risk construction activity with no assurance they will solve hot spot problems even at a cost much greater than remediation)

The changes outlined in this paper will yield major reliability and stability improvements while increasing useable raised floor space and the IT yield on site infrastructure investment. The many benefits include protecting IT availability while saving energy, reducing operating expenses, and reducing cooling unit capital investment.

As heat densities rise above 100 W/ft^2 (1,077 W/m^2) over large areas, other cooling technology solutions may be required. However, the most economical solution to high-density head loads may be to simply spread computer hardware equipment out, creating intentional "thermal footprint" white space, reducing the computer room average to what existing cooling technologies can easily manage. Another average reduction method is to partially fill racks, creating intentional vertical white space. When intentionally leaving rack unit spaces empty, filler plates must be used to separate the front from the rear of the rack or internal re-circulation of hot exhaust air can occur especially if cabling obstructions or inadequate rear door perforations exist.

Acknowledgements

The Uptime Institute, Inc. conducted the research outlined in this paper, under contract from Triton Technology Systems, Inc. The editorial content of this white paper is solely under the control of the *Institute*. Both companies are under common ownership. This white paper is used by both organizations.

Other Sources of Computer Room Cooling Information

The Uptime Institute has published numerous white papers on computer room cooling issues including:

- *2005-2010 Heat Density Trends in Data Processing Systems, and Telecommunications Equipment,*
- *Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms,*
- *Continuous Cooling Is Required for Continuous Availability, and*
- *Zinc Whiskers Growing on Raised Floor Tiles Are Causing Conductive Contamination Failures and Equipment Shutdowns.*

For a current listing of white paper information, go to www.uptimeinstitute.org/whitepapers.

About the Authors

Robert F. ("Dr. Bob") Sullivan, PhD, assisted in later stages by Lars Strong, P.E., developed the field air flow and cooling unit measurement methodology and performed the fieldwork for this two year study. W. Pitt Turner IV, P.E. provided technical peer review. The study was under the general supervision of Kenneth G. Brill, Executive Director of the *Institute*.

About the Uptime Institutue

The Uptime Institute, Inc. is a pioneer in creating and operating knowledge communities for improving uptime effectiveness in data center Facilities and Information Technology organizations. The 68 members of the *Institute's* Site Uptime® Network are committed to achieving the highest levels of availability with many being Fortune 100 companies. They interactively learn from each other as well as from *Institute* sponsored meetings, site tours, benchmarking, best practices, uptime effectiveness metrics, and abnormal incident collection and trend analysis. From this interaction and from client consulting work, the *Institute* prepares white papers documenting Best Practices for use by Network members and for the broader uninterruptible uptime industry. The *Institute* also conducts sponsored research and offers insightful seminars and training in site infrastructure management.

THE
UPTIME
INSTITUTE®

Building 100
2904 Rodeo Park Drive East • Santa Fe, NM 87505
Fax (505) 982-8484 • Phone (505) 986-3900
tui@uptimeinstitute.org • www.uptimeinstitute.org

TUI 811