



CONTINUOUS COOLING IS REQUIRED FOR CONTINUOUS AVAILABILITY:

Data Center Managers Need to Match Their Availability Expectations to the Appropriate Class of Continuous Cooling

By Rob Menuet, PE and W. Pitt Turner IV, PE

Technology compaction is producing ever-increasing power and heat densities in computer equipment racks and computer rooms. As densities continue to increase, any interruption in cooling will result in abrupt temperature increases far exceeding hardware manufacturer recommendations. This serious risk occurs because Uninterruptible Power Systems (UPS) use their batteries to continue providing uninterrupted power to computer and communication equipment during a utility power failure, while the cooling required for heat removal stops. As a result computer equipment intake temperatures must rise. Actually, this is only one of several common examples of how cooling can be halted while the generation of heat continues without interruption. In the past, actual heat densities were low enough that the thermal mass of the computer room air provided sufficient ride-through for short periods of lost cooling. This is no longer the case, and many managers will be perplexed by this “sudden” change in cooling requirements and how it affects their information availability and hardware reliability.

This white paper:

- Provides background on how inadequate cooling affects information reliability
- Gives Critical Physical Layer¹ managers the conceptual information they need to understand needed cooling changes
- Defines three classes of continuous cooling to assure abrupt temperature increases remain within manufacturer warranty limits
- Defines low, medium, high, and ultra-high heat density zones which are a function of both kW/rack and the quantity of such racks in a local area
- Helps managers select the appropriate class of continuous cooling, based on zone density and on the criticality of the Information Technology (IT) application/mission
- Provides guidance to design engineers in explaining, selecting, and implementing the selected continuous cooling class for both computer and UPS equipment rooms

Ever Increasing Product Heat Density

Chip, computer, communication, and storage manufacturers and proactive users have become increasingly concerned about rising power and heat densities. Getting more power to computers is not a difficult technical problem, but successfully removing the resulting heat is. If computer users cannot successfully and reliably cool these new devices, they will not be able to buy and deploy them. In 1997, 15 computer and communication hardware companies started cooperating on the development of a 10-year projection of future product heat densities. The Uptime Institute, Inc.[®] (*the Institute*) provided explanatory text and published the group’s projections in 2000 as *2000-2010 Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment*². Figure 1 re-prints the *2000-2010 Product Heat Density Projections* per product footprint. It is important to note that the average heat density across the entire computer room is substantially lower than the product footprint figures due to the space occupied by power and cooling equipment, required product service access and service aisles, code required safety spacing, columns, and other “gross” space consumers which substantially reduce the “net” space available for the computer and communication hardware.³

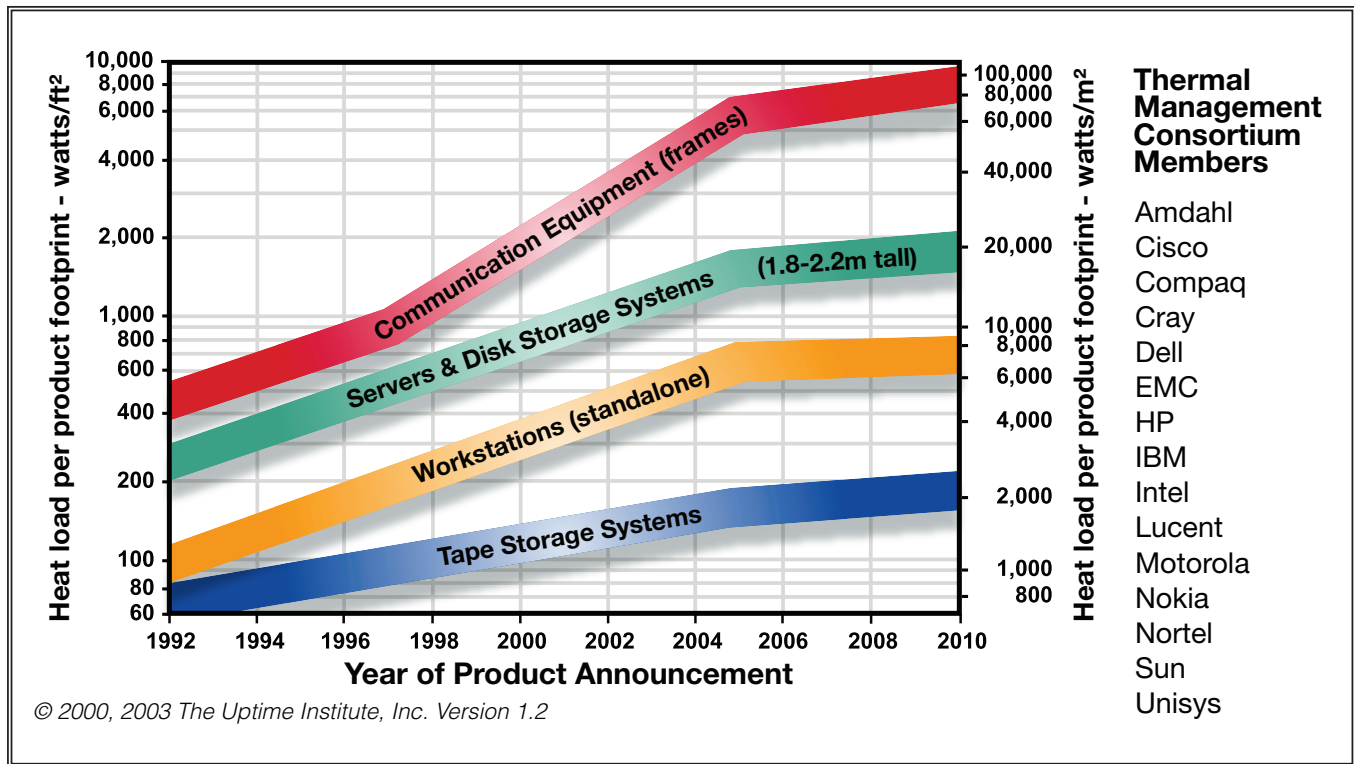
¹ The *Institute* defines the Critical Physical Layer as including the physical housing and installation of computer and communication hardware plus the necessary computer room cabinets, space, power, cooling, cabling, and other critical environment systems required for the physical hardware to operate reliably. The Critical Physical Layer is the missing foundation layer in the seven-level Open Systems Interface protocol.

² The 2005-2010 trends booklet is available at www.uptimeinstitute.org/whitepapers.

³ To convert from product footprint density to the gross computer room average, divide the product footprint figure by a minimum of 5 which corresponds to 30 ft² (2.8 m²) of gross computer room space per rack or cabinet. Most computer rooms fail to achieve this level of packing efficiency. For such rooms, use a divisor greater than 5.



**Figure 1:
2000-2010 Product Heat Density Projections**



It is important to note that the scale on the Y axis in Figure 1 is logarithmic. If these trends had been plotted to a linear scale they would have abruptly risen vertically off the chart. The slope of the Server and Disk Storage Systems band is a compounded annual rate of growth of 17%. The steep growth rate in this chart was greeted with widespread industry skepticism when it was initially released.

The 2000 projections were updated in 2005 as it had become apparent during the intervening time that the original product projections had actually been *far too conservative*. Actual product releases had significantly exceeded the predicted trend line by 2003. Figure 2 presents the updated 2005 projections along with actual datapoints for several products released in 2004 and 2005. (These datapoints are based on actual product power consumption.)

The 2000 and 2005 projections involve several important assumptions. The first is what the heat densities will be for next, next, next generation products which aren't even in design today. The second is an assumption that every available opening will be occupied with equipment.

In the real world, manufacturers do not control how customers actually lay out their rooms, nor how they fill their racks and cabinets with equipment. Finally, the projections assume 100% new equipment whereas actual computer rooms are a "technology stew" of old-old, old, and the very latest technology. In addition, some applications do not require high-power computing whereas others do. This results in an "application stew" that requires different types of processing performance. These "stew" factors turn out to be critical variables in understanding and predicting actual heat densities.

Actual Computer Room Heat Densities Have Grown, But Not at Predicted Rates

Actual computer room heat densities have grown, but not at the same rate as projected by the manufacturers. The *Institute* has done benchmarking of gross computer room densities for a group of nineteen *Site Uptime Network*⁴ member sites comprising a total of almost 1,000,000 ft² (93,000 m²) of computer room floor space.

⁴The *Site Uptime Network (Network)* is a closed learning community of 68 mostly large organizations, who do benchmarking and other activities to improve their site uptime. The *Network* is organized and facilitated by the *Institute*.



Consistent, validated, annual data available from Q4 of 1999 to Q4 of 2005 shows an average heat density of 32 W/ft² (345 W/m²) versus 23 W/ft² (248 W/m²) for 1999. This is an increase of 39% over 6 years or an annual growth of 6.5 % per year. (Since these are average figures calculated over very large areas, localized spot and zone densities were certainly a great deal higher, and were most likely similar to those shown in the density projections.)

The computer room with the highest average density was 65W/ft² (700 W/m²) in 2005 versus 23 W/ft² (248 W/m²) in 1999. In 2005, eight rooms or 42% of the total group had densities equal to or exceeding 40 W/ft² (431 W/m²) and 17 rooms or 90% of the 1,000,000 ft² exceeded 20W/ft² (1.9W/m²). For discussion on actual site data and possible reasons for the large variances between the projected and actual heat densities, please see the *Institute's* 2006 version of *2005-2010 Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment*. For the purposes of this white paper, the relevant point is that real world heat densities have grown significantly since 1999 with 90% of the rooms studied exceeding 20 W/ft² (215 W/m²) and

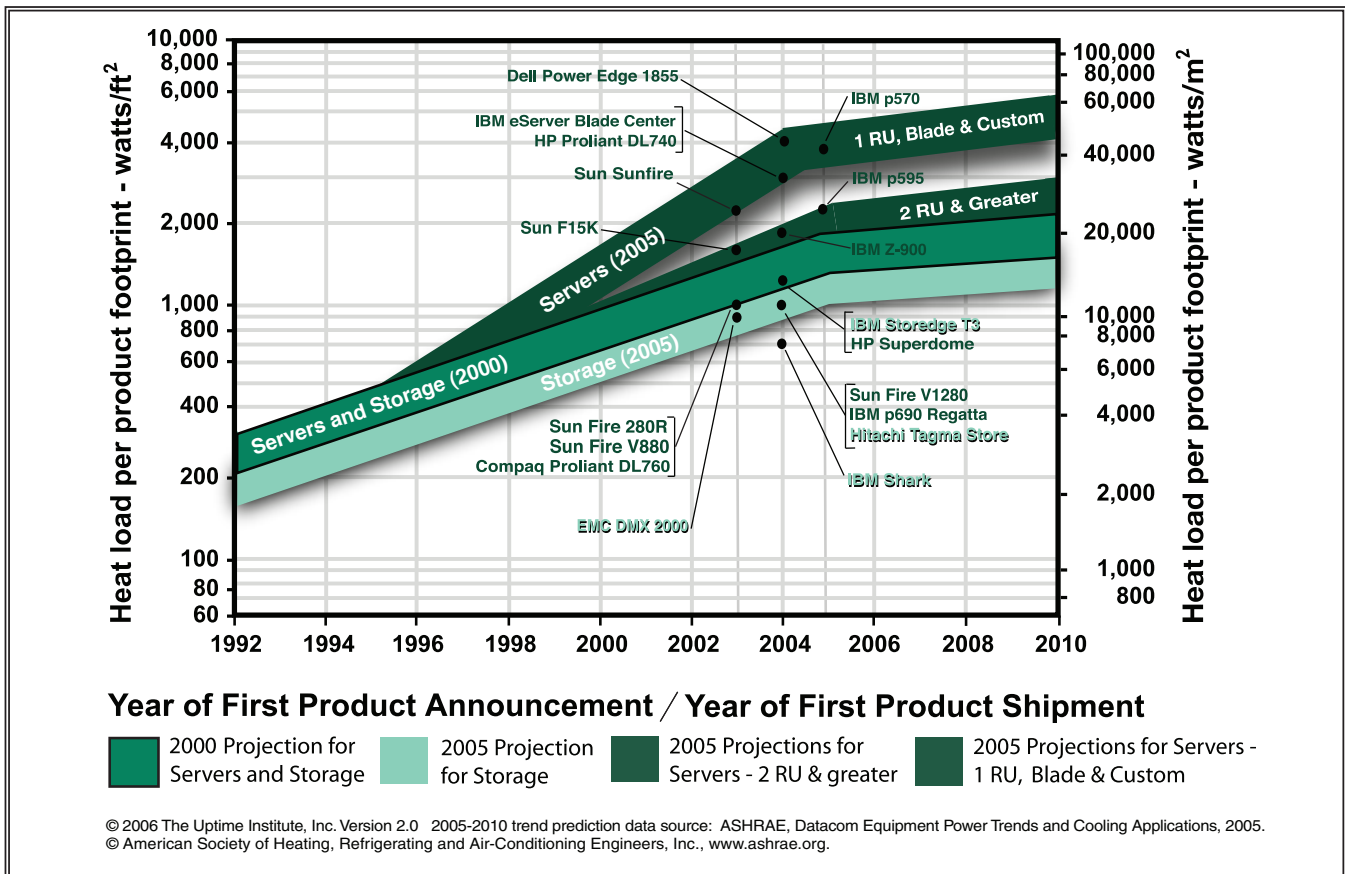
42% exceeding 40 W/ft² (430 W/m²). The significance of these two zone density benchmarks will become apparent as data is presented on temperature rise upon loss of cooling. However, before getting to that, the question of whether temperature and relative humidity are important to information availability must be addressed.

Temperature as a 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

**Figure 2:
2005-2010 Product Heat Density Projections**





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 3 which follows along with the *Institute's* recommended cooling unit return air control temperature and relative humidity set-points.

A number of very important points can be drawn from Figure 3

- 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% Rh). This allows for optimal cooling unit efficiency with some tolerance for local hot spots.
- The Institute has done extensive research on actual cooling conditions in 19 computer rooms comprising 204,000 ft² (19,000 m²). This research found that the actual air intake temperatures exceeded guidelines

for 10% of the racks. On average, only 40% of the available supply of cold air was directly cooling computer equipment. The other 60% was bypassing the computer equipment dramatically reducing cooling effectiveness. Within the study, the rooms with the highest excess cooling capacity also had the greatest number of hot spots. For example, the room experiencing 28% hot spots also had 5.2 times more cooling capacity running than was required by the heat load. (This data indicates many critical layer physical managers must first learn to successfully master the art of cooling low densities before they can begin implementing medium, high, or ultra-high densities.)⁵

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

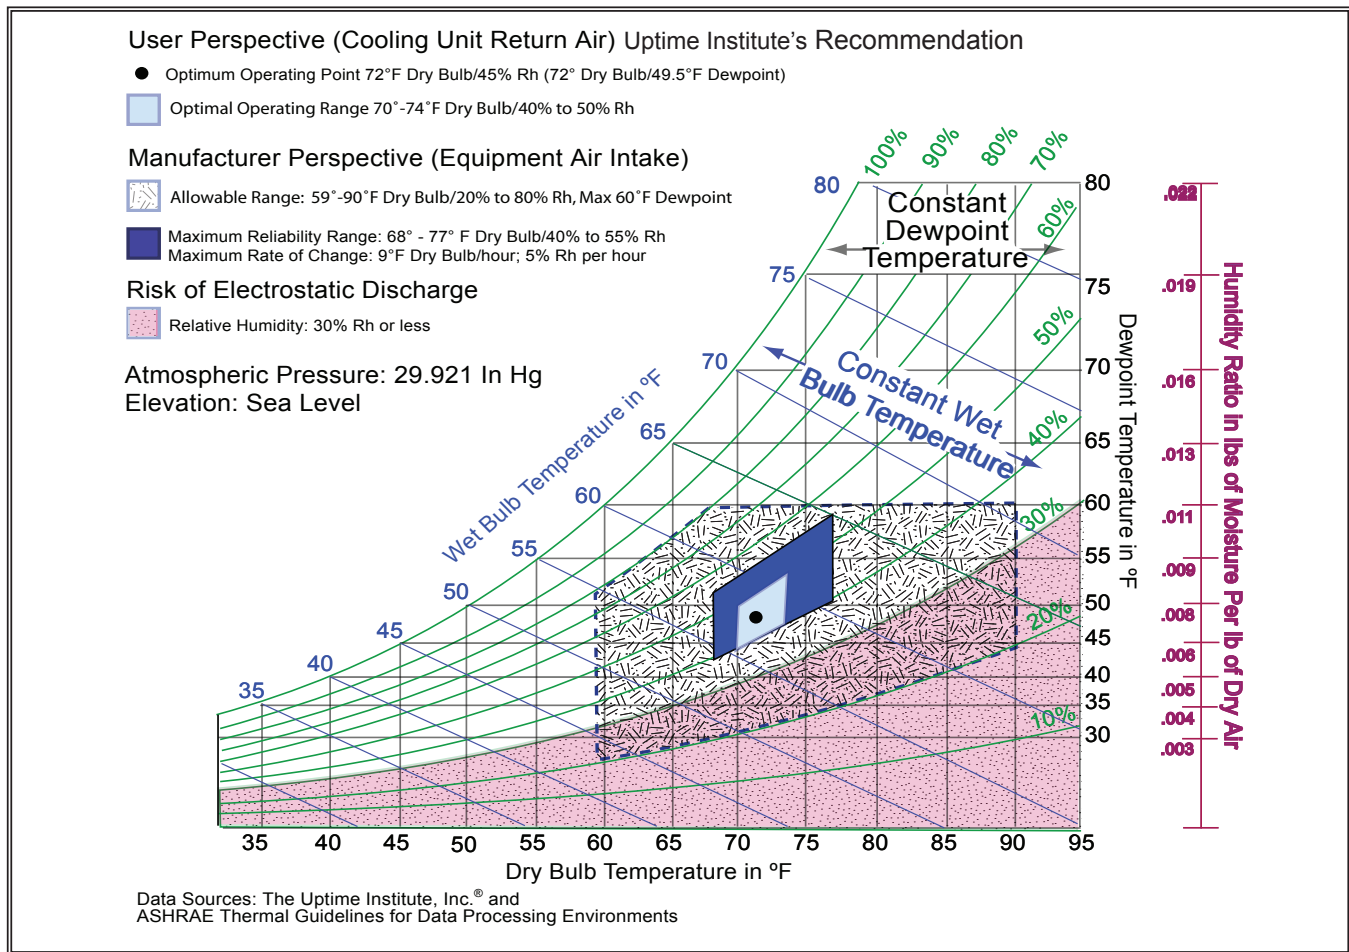
Excessive expansion and contraction of the component materials inside computer and communication equipment results from exceeding manufacturer temperature range and rate of change recommendations. When the equipment in a computer room is "cooked," a considerably higher rate of premature equipment failure 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 cycle⁶ is to re-seat all cards

⁵ For further information, see *Reducing Bypass Airflow Is Essential For Eliminating Computer Room Hot Spots* available at www.uptimeinstitute.org/whitepapers

⁶ 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

**Figure 3:
Computer Hardware Reliability Environmental Reliability Limits**



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 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 6 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



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.

How Do Excessive Thermal Excursions Occur?

Escalating computing equipment heat densities resulting from normal technology compaction trends has created conditions where intake air temperatures in computer rooms will quickly soar to damaging levels within 60 seconds (depending upon density) after a loss of cooling air movement. Making matters worse, additional damage is possible after cooling is restored and the room is cooled down too quickly. These issues have always been present; what is new and unprecedented is the rate at which temperatures will rise after a cooling failure because of increased densities.

One common example of a cooling loss is a utility failure where the UPS continues pumping energy into the computer room which the hardware converts to heat. Even if an engine generator is available, starts, and successfully assumes the load⁷, there can easily be several minutes during which there is no cooling because the cooling equipment is being cycled back online. If this heat is not continuously removed, computer room temperatures must rise (think of an enclosed box the size of a refrigerator containing 300 100-W light bulbs and what would happen with no cold air circulation).

Another common cause of cooling loss occurs when the cooling equipment (a chiller, for example) locks itself out for any number of reasons. It can easily take 15 to 30 minutes for cooling to be restored. In the middle of the night or on weekends, when no facility staff is on site to manually override safety circuits, it can take even longer to get cooling back online.

How High and How Fast Can Temperatures Rise After a Loss of Cooling?


ComputerSite Engineering, Inc. (a sister company of the *Institute*) has conducted extensive field testing on what happens to ambient temperatures when cooling is intentionally interrupted in different ways and under different load conditions. Because of the serious threat to hardware reliability, such tests can only be conducted using load banks that simulate the heat produced by computer equipment when it is powered up. This results of these extensive field tests collected in a 30,000 ft² (2,800 m²) site is shown in Figure 4.

The most important conclusions are that average room temperatures in hot and cold aisles after cooling was interrupted rose by:

- 16°F (9°C) in 14 minutes to 87.5°F (30.8°C) with a historically modest heat load of only 20 W/ft² (215 W/m²) or (0.6 kW/rack). (Seventeen out of 19 rooms exceeded this density in Q4 of 2005.)
- 25°F (14°C) in 10 minutes to 96°F (35.5°C) at what today is only a modest density of 40 W/ft² (431 W/m²) or 1.2 kW/rack). (Eight out of the 19 previously mentioned computer rooms had a Q4 of 2005 density exceeding 40 W/ft² (431 W/m²) and had concentrations far exceeding this in localized areas.)
- Since total moisture in the room remained constant, relative humidity dropped to 25% Rh and 19% Rh respectively, or well down into the electrostatic discharge region.
- These are average conditions, local conditions in areas of high equipment density concentrations would have been considerably hotter and the Rh even lower.

⁷ Site Uptime Network data indicates that at least once a year the cooling fails at one member site for an extended period despite the engine generators starting successfully because they didn't successfully transfer load until someone manually intervened. Typically, it takes almost the full duration of the UPS battery discharge (typically 15 minutes) to figure out what is wrong and to manually override the problem.

Figure 4:
Ambient Computer Room Temperature Rise during a Cooling Failure while the UPS Continues to Supply Critical Power.



These environmental conditions far exceed manufacturer recommendations for maximum reliability and performance. Just to repeat, these intake conditions are:

- Dry bulb temperature less than or equal to 77°F (25°C) and rate of change of less than 9°F (5°C) per hour
- Relative humidity equal to or greater than 40% Rh and a rate of change less than 5% per hour

In addition to totally interrupting all cooling, experiments were run on just interrupting the production of new cooling while still maintaining air circulation within the computer room. Maintenance of air circulation keeps the air in the computer room stirred up and it turns out there is a very significant heat transfer with the cold soaked skin of the room. Instead of the abrupt temperature increases experienced when all cooling was stopped, maintaining cooling unit air circulation resulted in a very gradual 4°F (2.2°C) temperature increase over 14 minutes when the experiment was stopped. The data trend would seem to indicate that the temperature would have continued to increase gradually if the experiment had been allowed to continue past 14 minutes.

A final, very important, lesson can be drawn from these experiments. Note the 23°F (13°C) temperature drop on Figure 4 in just 4 minutes (from T= 10:00 to T= 14:00) after the experiment ended and cooling was restored. This rate of change far exceeds manufacturer recommendations and should be avoided in a real world environment. The natural thing will be to restore full cooling, and special technician training will be required to overcome this instinctive urge. Without rehearsal and specific procedures, controlling the rate of temperature drop in the pandemonium of restoring availability is likely to be overlooked or disregarded.

Density Increases Tremendously Upon Cooling Failure

Even in rooms having low overall equipment densities, local areas will have much higher equipment density concentrations. Such concentrations can be either spot (one or two racks) or across entire equipment zones (a building bay). Such localized concentrations are normal and almost unavoidable.



However, these local concentrations are further increased upon loss of cooling air flow because the effective gross computer room area is reduced while the power consumption remains the same.

Upon loss of cooling, the only air movement will be from the internal equipment fans which are continuing to receive uninterrupted power. These fans will force hot air exhaust into the Hot Aisle while creating suction at the equipment air intake. This suction will tend to recirculate the exhausted hot air back into the computer equipment air intake. While some mixing from adjacent racks will occur, the effective local area is approximately 14 to 16 ft² (1.3 to 1.5m²) which is composed of the cabinet width of approximately 2 ft (.6m) and half the repeating pitch of the Cold Aisle/ Hot Aisle or between 7 and 8 ft. (2.1 to 2.4m). (This is less than half the normal area when cooling equipment is keeping the ambient air stirred so the resulting local density can be predicted to more than double.)

Spot concentrations involving one or two racks will permit some mixing of the hot exhaust air with adjacent stagnant cooler air which will average down the resulting temperature rise. Within a zone of high density equipment, there can be little mixing because everything being circulated is hot.

When the actual kW/rack density in a local area or zone exceeds the room average, the local recirculation pattern combined with a higher heat load can result in densities exceeding 300 W/ft² (3,231 W/m²), which takes only 4.5 kW per rack of power dissipation.⁸ At this density, theoretical computational fluid dynamics modeling shows that a critical temperature rise will occur in less than one minute. Under such circumstances, many servers will thermal themselves off as a protective measure.

UPS Equipment Rooms Must Also Be Considered

UPS systems utilize solid state circuitry and are susceptible to the same high temperature concerns as the computer and communication equipment. It is important to provide continuous UPS equipment room cooling so that during a UPS discharge event, the room temperature does not increase to the point where the stability and reliability of the UPS is affected.

Continuous Cooling Classes Defined

As shown in Figure 4, the rate of temperature increase upon loss of cooling can be greatly diminished if air circulation of the computer room cooling units is maintained. This can be accomplished by putting fans in the cooling units on a UPS power source.

To codify the lessons learned, the *Institute* has created three classes of continuous cooling which are defined as

- **Class A Uninterruptible Cooling:** Computer room cooling unit fans and secondary chilled water pumps are on UPS power with a chilled water thermal storage tank. The objective is to keep the entire computer room at ‘normal’ temperature and relative humidity conditions during an event that would otherwise halt cooling.
- **Class B Continuous Cooling:** Computer room cooling unit fans are on UPS, but circulation pumps are not. The objective is to keep the computer room air circulating during this event.
- **Class C Interruptible Cooling:** No cooling equipment is on UPS. This is commonly seen in many current designs which subject the computer room to loss of cooling during a power failure or during thermal cycling after restoration of power.

⁸ As a point of reference, fully stacked blade and 1 RU servers already generate between 8 and 14 kW per rack, and 30kW to 50 kW per rack products are on hardware manufacturer drawing boards for near term release. Using this same short cycling analysis with these hotter next generation products, unacceptable temperature rises will occur practically instantaneously.



Cautions on achieving continuous cooling

- Selection of the UPS source of power for cooling unit fans and circulation pumps should be carefully studied. Different design options include using the same UPS as the computer equipment loads, using the same UPS as used for computer loads but using isolation transformers for the mechanical loads, or using a separate and independent UPS system. Owners and their design professionals should understand the risks and benefits of each option before making a selection.
- Designs which transfer mechanical loads to a UPS only after a utility power failure can be very risky and need careful thought to prevent cascading failures. Fans and pumps selected for UPS power should be on that source for their normal operation.
- For Class A Uninterrupted Cooling, the duration of thermal storage should at a minimum match the discharge duration of the UPS battery plant. Additional thermal storage can be used to protect the critical load from chiller plant or automation issues by providing additional time for diagnosis and solution of issues preventing automatic restoration of cooling production. Different locations for the thermal storage tank in the chilled water piping system can dramatically affect the recovery time to restore optimal computer room temperature after a disruption.

Mechanical Cooling System Types

Some mechanical system types are more easily adaptable to Class A Uninterruptible Cooling than others. This section addresses how to adapt the more common types of mechanical systems to the three classes of continuous cooling.

Air-cooled direct expansion (DX) mechanical systems which employ compressors in each computer room air conditioning (CRAC) unit pose a significant design and cost challenge for achieving Class A Uninterruptible Cooling, as it is generally impractical and expensive to power compressors through a UPS system. For air-cooled DX CRAC type mechanical systems, Class B or C cooling may

be the only feasible choices. For Class B Continuous Cooling, the cooling units must be configured with a separate power source for the fan/control sections than for the compressors. This will allow the air circulation fans to be powered through the UPS without the additional burden of compressor loads allowing the fans to continue operating during the power outage.

Water and glycol cooled DX systems with free cooling cycles may be configured to provide Class A Uninterruptible Cooling, although it is very complex, awkward, and not generally recommended. If DX is chosen, the design method is to create and store chilled water as a thermal “battery,” which can then be circulated in the event of a power outage. The ride-through duration of time the system can function is limited only by the size of the chilled water storage tank. The economizer control sequences, inherent with the CRAC, unit senses the chilled water entering the cooling unit, and diverts the chilled water to the economizer coils, thereby offering an active cooling source without the need to operate the compressors. In this scenario, a separate chiller plant will be required to maintain the chilled water source during normal operating periods, but this plant does not require UPS power. Again, in this scenario, the CRAC units should be configured with separate power supplies for the fan/control sections and the compressors.

Chilled water systems offer the best opportunity for Class A Uninterruptible Cooling. Since the computer room air-handling (CRAH) units do not have compressors, separating the blower and compressor power supplies is not necessary. Stored chilled water is readily created by inserting a chilled water storage tank in the appropriate location in the chilled water loop to provide the thermal reserve. For this type of mechanical system, CRAH unit fans and the chilled water circulating pumps are provided with UPS power.



Whenever thermal storage is provided, it is likely to be very economical to incrementally increase the size of thermal storage to achieve other operational advantages such as utility peak shaving and protection from short term chiller lock outs or building automation problems. Sites which have thermal storage are generally surprised by how frequently they are used in day-to-day operation and facility managers wish their sizes were much larger.

Defining Low, Medium, High, and Ultra-High Density Zones

Over the last several years, the *Institute* has been urging its member companies to begin using kilowatts (kW) per rack or cabinet instead of watts per square foot (square meter). Kilowatts per rack or cabinet is much less ambiguous than watts per unit of area, which is subject to misinterpretations as to what is included within the area. Manufacturers will focus on product footprint, IT will tend to focus on product footprint plus ½ service access, facilities will focus on gross computer room, and the electric utility account executive will focus on the total building footprint. The potential for confusion is tremendous, whereas kW per rack or cabinet is unambiguous.

Using actual kW per rack or cabinet as metric, the *Institute* uses the following definitions for the average density in a zone which includes 33 racks or cabinets (approximately 1,000 ft² (93 m²):

- Low density < 0.6 kW per rack or cabinet (<20 W/ft² or 215 W/m²)
- Medium density >0.6 and < 1.2 kW per rack or cabinet (<40 W/ft² or 430 W/m²)
- High density > 1.2 kW < 3.3 kW per rack or cabinet (<110W/ft² or 1,184 W/m²)
- Ultra-high density > 4.0 kW per rack or cabinet (>133 W/ft² or 1,432 W/m²)

As a useful point of reference, only two out of the 19 computer rooms studied at the end of 2005 had actual loads consistent with low density (i.e., the actual loads in 17 out of 19 exceeded low density) and 42 % (i.e., 8 out of 19) exceeded the threshold for medium density.

Selecting the Appropriate Class of Cooling Based on Zone Density and IT Mission

The choice of the appropriate class of continuous cooling is a function of both zone density and the required information availability. A lower density zone will not reach damaging temperatures as quickly. As spot and zone densities increase, the need for Class A Uninterruptible Cooling increases. Sites with mission critical high availability applications should consider Class A Uninterruptible Cooling regardless of their density.

The *Institute's* guidance is interruptible Class C cooling is most appropriate for low density zones, continuous Class B cooling is recommended for medium density and high density, and uninterruptible Class A cooling is required for ultra-high density.

Within a Class C cooling zone, load in several of the racks can greatly exceed the zone average, provided they are not immediately adjacent or near to each other. The separation required will be a function of room configuration and the relative loads. Enough spacing must be provided to allow the re-circulated hot air mixing driven by the internal hardware fans in the higher kW racks to average down the spot density.

Despite having high or ultra-high densities, some customers and some applications will not be able to justify the added cost of Class A Uninterruptible Cooling. Critical Layer Managers may choose to accept the risk of equipment damage or unreliability. However, this could be a career limiting decision. The world has abruptly changed in a non-linear manner, and managers who predict the future by looking at the past may not anticipate the profound impact that the foundational changes discussed in this paper will have on their business.

Hardware manufacturers now include imbedded temperature sensors in all high end products



which report actual internal temperatures back to their service organizations. The customer is no longer the only one who knows that a high temperature excursion has occurred. This information is now available to manufacturers in an inarguable format. This information over the longer term is likely to affect their decisions about honoring warranty contracts, which in the past have transferred customer environmental problems to manufacturers.

Conclusion

The effects of technology compaction are producing ever-increasing densities. Designers and managers of high power density computer environments must provide continuous cooling to prevent unexpected thermal damage to the computing equipment. The decision on which class of continuous cooling is required is a function of both the density class and IT availability requirements.

About the Authors

Mr. Menuet is a Senior Principal of GHT Limited in Arlington, VA. He is a frequent technical contributor to The Uptime Institute.

Mr. Turner is a Distinguished Fellow and Senior Certification Authority for the *Institute* and a Principal of ComputerSite Engineering, Inc. in Santa Fe, NM.

Additional editorial guidance has been provided by Kenneth G. Brill and Robert “Dr. Bob” F. Sullivan, PhD. Mr. Brill is the founder of the *Institute* and is its Executive Director. Dr. Sullivan performed much of the *Institute’s* original computer room cooling research.

About The Uptime Institute

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.

© 2006 The Uptime Institute, Inc.



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