© 2004. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. (www.ashrae.org). Reprinted by permission from ASHRAE Transactions, Vol. 110, Part 1. This material may not be copied nor distributed in either paper or digital form without ASHRAE's permission.

AN-04-9-1

Evolution of Data Center Environmental Guidelines

Roger R. Schmidt, Ph.D.
Member ASHRAE

Christian Belady
Associate Member ASHRAE

Alan Classen

Tom Davidson
Member ASHRAE

Magnus K. Herrlin, Ph.D.
Member ASHRAE

Shlomo Novotny

Rebecca Perry

ABSTRACT

Recent trends toward increased equipment power density in data centers can result in significant thermal stress, with the undesirable side effects of decreased equipment availability, wasted floor space, and inefficient cooling system operation. In response to these concerns, manufacturers identified the need to provide standardization across the industry, and in 1998 a Thermal Management Consortium was formed. This was followed in 2002 by the creation of a new ASHRAE Technical Group to help bridge the gap between equipment manufacturers and facility designers and operators. "Thermal Guidelines for Data Processing Environments" the first publication of TC9.9, is discussed in this paper, along with a historical perspective leading up to the publication and discussion of issues that will define the roadmap for future ASHRAE activities in this field.

CURRENT INDUSTRY TRENDS/PROBLEMS/ISSUES

Over the years, computer performance has significantly increased but unfortunately with the undesirable side effect of higher power. Figure 1 shows the National/International Technology Roadmap for Semiconductors' projection for processor chip power. Note that between the years 2000 and 2005 the total power of the chip is expected to increase 60% and the heat flux will more than double during this same period. This is only part of the total power dissipation, which increases geometrically. The new system designs, which include very efficient interconnects and high-performance data-bus design, create a significant increase in memory and other device utilization, thus dramatically exceeding power dissipation expectations. As a result, significantly more emphasis has been placed on the cooling designs and power delivery methods within electronic systems over the past year.

In addition, the new trend of low-end and high-end system miniaturization, dense packing within racks, and the increase in power needed for power conversion on system boards have caused an order of magnitude rack power increase. Similarly, this miniaturization and increase in power of electronics scales into the data center environment. In fact, it wasn't until recently that the industry has publicly recognized that the increasing density within the data center may have profound impact on the reliability and performance of the equipment it houses in the future. For this reason, there has been a recent flurry of papers addressing the need for new room cooling

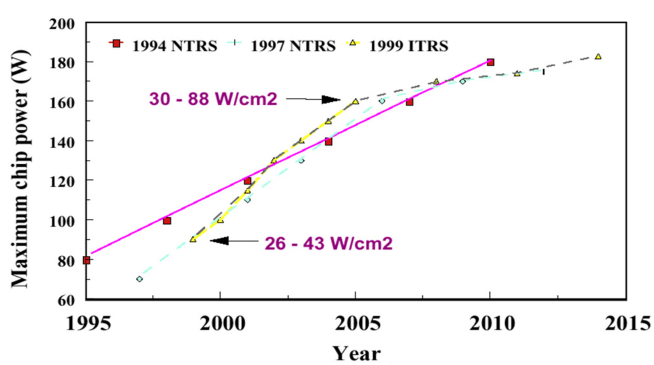

Figure 1 *Projection of processor power by the National/International Technology Roadmap for Semiconductors.*

Roger Schmidt and **Alan Claassen** are with IBM Corp., San Jose, Calif. **Christian Belady** is with Hewlett-Packard, Richardson, Tex. **Tom Davidson** is with DLB Associates, Ocean N.J. **Magnus Herrlin** is with ANCIS Incorporated, San Francisco, Calif.
Shlomo Novotny and **Rebecca Perry** are with Sun Microsystems, San Diego, Calif.

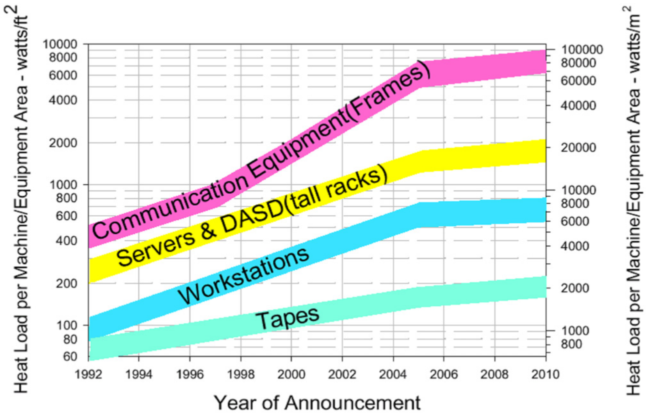

Figure 2 *Equipment power projection (Uptime Institute).*

technologies as well as modeling and testing techniques within the data center. All of these recognize that the status quo will no longer be adequate in the future. So what are the resulting problems in the data center? Although there are many, the following list discusses some of the more relevant problems:

1. Power density is projected to go up

   Figure 2 shows how rapidly machine power density is expected to increase in the next decade. Based on this figure it can easily be projected that by the year 2010 server power densities will be on the order of 20,000 W/m^2. This exceeds what today's room cooling infrastructure can handle.

2. Rapidly changing business demands

   Rapidly changing business demands are forcing IT managers to deploy equipment quickly. Their goal is to roll equipment in and power on equipment immediately. This means that there will be zero time for site preparation, which implies predictable system requirements (i.e., "plug and play" servers).

3. Infrastructure costs are rising

   The cost of the data center infrastructure is rising rapidly with current costs in excess of about \$1000/ft^2. For this reason, IT and facility managers want to obtain the most from their data center and maximize the utilization of their infrastructure. Unfortunately, there are many barriers to achieve this.

   First, airflow in the data center is often completely ad hoc. In the past, manufacturers of servers have not paid much attention to where the exhaust and inlets are in their equipment. This has resulted in situations where one server may exhaust hot air into the inlet of another server (sometimes in the same rack). In these cases, the data center needs to be overcooled to compensate for this inefficiency.

   In addition, a review of the industry shows that the environmental requirements of most servers from various manufacturers are all different, yet they all coexist in the same environment. As a result, the capacity of the data center needs to be designed for the worst-case server with the tightest requirements. Once again, the data center needs to be overcooled to maintain a problematic server within its operating range.

   Finally, data center managers want to install as many servers as possible into their facility to get as much production as possible per square foot. In order to do this they need to optimize their layout in a way that provides the maximum density for their infrastructure.

   The above cases illustrate situations that require overcapacity to compensate for inefficiencies.

4. There is no NEBS equivalent specification for data centers. (NEBS [Network Equipment–Building Systems] is the telecommunication industry's most adhered to set of physical, environmental, and electrical standards and requirements for a central office of a local exchange carrier.)IT/facility managers have no common specification that drives them to speak the same language and design to a common interface document.

The purpose of this paper is to review what started as a "grassroots" industrywide effort that tried to address the above problems and later evolved into an ASHRAE Technical Committee. This committee then developed "Thermal Guidelines for Data Processing Environments" (ASHRAE 2003a), which will be reviewed in this paper.

HISTORY OF INDUSTRY SPECIFICATIONS

Manufacturers Environmental Specifications

In the late 1970s and early 1980s, data center site planning consisted mainly of determining if power was clean (not connected to the elevator or coffee pot), had an isolated ground, and if it would be uninterrupted should the facility experience a main power failure. The technology of the power to the equipment was considered the problem to be solved, not the power density. Other issues concerned the types of plugs, which varied widely for some of the larger computers.

In some cases, cooling was considered a problem and, in some isolated cases, it was addressed in a totally different manner, so that the technology and architecture of the machine were dictated by the cooling methodology. Cray Research, for example, utilized a liquid-cooling methodology that forced a completely different paradigm for installation and servicing. Large cooling towers, which were located next to the main system, became the hallmark of computing prowess. However, typically the preferred cooling methods were simply bigger, noisier fans. The problem here was the noise and the hot-air recirculation when a system was placed too close to a wall.

Over the last ten years, the type of site planning information provided has varied depending on the company's main product line. For companies with smaller personal computers or workstations, environmental specifications were much like those of an appliance: not much more than what is on the back of a blender. For large systems, worksheets for performing calculations have been provided, as the different configurations had a tremendous variation in power and cooling requirements. In certain countries, power costs were a huge contributor to total cost of ownership. Therefore, granularity, the ability to provide only the amount of power required for a given configuration, became key for large systems.

In the late 1990s, power density became more of an issue and simple calculations could no longer ensure adequate cooling. Although cooling is a factor of power, this does not provide the important details of the air flow pattern and how the heat would be removed from the equipment. However, this information was vitally needed in equipment site planning guides. This led to the need for additional information, such as flow resistance, pressure drop, and velocity, to be available not only in the design stages of the equipment but after the release of the equipment for integration and support. This evolved into the addition of more complex calculations of the power specifications, plus the airflow rates and locations, and improved system layouts for increased cooling capacity.

In the early 2000s, power densities continued to increase as projected. Layout, based on basic assumptions, could not possibly create the efficiencies in the airflow that were required to combat the chips that were scheduled to be introduced in the 2004 time frame. Because of this, thermal modeling, which was typically used to design products, began to be viewed as a possible answer for optimizing cooling in a data center environment. However, without well-designed thermal models from the equipment manufacturers and easy-to-use thermal modeling tools for facility managers, the creation or troubleshooting of a data center environment fell again to traditional tools for basic calculations to build a data center or to gather temperatures after a data center was in use.

At this point it became apparent that the solution for rising heat densities could not be solved after the delivery of a product. Nor could it be designed out of a product during development. Instead, it had to be part of the architecture of the entire industry's next-generation product offerings.

Formation of Thermal Management Consortium

In 1998 a number of equipment manufacturers decided to form a consortium to address common issues related to thermal management of data centers and telecommunications rooms. Initial interest was expressed from the following companies: Amdahl, Cisco Systems, Compaq, Cray, Inc., Dell Computer, EMC, HP, IBM, Intel, Lucent Technologies, Motorola, Nokia, Nortel Networks, Sun Microsystems, and Unisys. They formed the Thermal Management Consortium for Data Centers and Telecommunications Rooms. Since the industry was facing increasing power trends, it was decided that the first priority was to develop and then publish a trend chart on power density of the industry's equipment that would aid customers in planning data centers for the future. Figure 2 shows the chart that resulted from this effort. This chart has been widely referenced and was published in collaboration with the Uptime Institute in 2000. Following this publication the consortium formed three subgroups to address what customers felt was needed to align the industry:

A. Rack airflow direction/rack chilled airflow requirements

B. Reporting of accurate equipment heat loads

C. Common environmental specifications

The three subgroups worked on the development of guidelines to address these issues until an ASHRAE committee was formed in 2002 that continued this effort. The result of these efforts is "Thermal Guidelines for Data Processing Environments" (ASHRAE 2003a), which is being published by ASHRAE. The objectives of the ASHRAE committee are to develop consensus documents that will provide environmental trends for the industry and guidance in planning for future data centers as they relate to environmental issues.

Formation of ASHRAE Group

The responsible committee for data center cooling within ASHRAE has historically been TC9.2, Industrial Air Conditioning. The *2003 ASHRAE Handbook—HVAC Applications*, Chapter 17, "Data Processing and Electronic Office Areas" (ASHRAE 2003b) has been the primary venue within ASHRAE for providing this information to the HVAC industry. There is also *Standard 127-2001, Method of Testing for Rating Computer and Data Processing Room Unitary Air-Conditioners* (ASHRAE 2001), which has application to data center environments.

Since TC9.2 encompasses a very broad range of industrial air-conditioning environments, ASHRAE was approached in January 2002 with the concept of creating an independent committee to specifically address high-density electronic heat loads. The proposal was accepted by ASHRAE, and TG9.HDEC, High Density Electronic Equipment Facility Cooling, was created. TG9.HDEC's organizational meeting was held at the ASHRAE Annual Meeting in June 2002 (Hawaii). TG9.HDEC has since further evolved and is now TC9.9, Mission Critical Facilities, Technology Spaces, and Electronic Equipment.

The first priority of TC9.9 was to create a thermal guidelines document that would help to align the designs of equipment manufacturers and aid data center facility designers to create efficient and fault tolerant operation within the data center. The resulting document, "Thermal Guidelines for Data Processing Environments," was completed in a draft version on June 2, 2003. It was subsequently reviewed by several dozen companies representing computer manufacturers, facilities design consultants, and facility managers. Approval to submit the document to ASHRAE's Special Publications Section was made by TC9.9 on June 30, 2003, and the document publication is expected in December 2003.

TC9.9 ENVIRONMENTAL GUIDELINES

Environmental Specifications

For data centers, the primary thermal management focus is on the assurance that the housed equipment's temperature and humidity requirements are being met. As an example, one large computer manufacturer has a 42U rack with front-to-rear air cooling and requires that the inlet air temperature into the front of the rack be maintained between 10°C and 32°C for elevations up to 1287 m (4250 feet). Higher elevations require a derating of the maximum dry-bulb temperature of 1°C for every 218 m (720 feet) above 1287 m (4250 feet) up to 3028 m (10000 feet). These temperature requirements are to be maintained over the entire front of the 2 m height of the rack where air is drawn into the system. These requirements can be a challenge for customers of such equipment, especially when there may be many of these racks arranged in close proximity to each other and each dissipating powers up to 12.5 kW when fully configured.

As noted in the example above, data processing manufacturers typically publish environmental specifications for the equipment they manufacture. The problem with these specifications is that other manufacturers with the same type of equipment and selling into the same customer environment may have a different set of environmental specifications. Not only do discrepancies occur between manufacturers; in some cases, there are discrepancies within the portfolio of a manufacturer's products. As one can imagine, customers of such equipment can be left in quandary as to what environment to provide in their data processing room.

In an effort to standardize the environmental specifications, the ASHRAE TC9.9 committee first surveyed the environmental specifications of a number of data processing equipment manufacturers. From this survey, four classes were developed that would encompass most of the specifications. Also included within the guidelines was a comparison to the NEBS specifications for the telecommunications industry to show both the differences and also aid in possible convergence of the specifications in the future.

The four data processing classes cover the entire environmental range from air-conditioned server and storage environments of classes 1 and 2 to the lesser controlled environments such as class 3 for workstations, PCs, and portables or class 4 for point-of-sales equipment with virtually no environmental control. For each class the allowable dry-bulb temperature, relative humidity, maximum dew point, maximum elevation, and maximum rate of change are specified for product operating conditions. For higher altitudes, a derating algorithm is provided that accounts for diminished cooling. In addition to the allowable ranges, the recommended range for dry-bulb and relative humidity is provided for classes 1 and 2 based on the reliability aspects of the electronic hardware. Non-operating specifications of dry-bulb, relative humidity, and maximum dew point are also included.

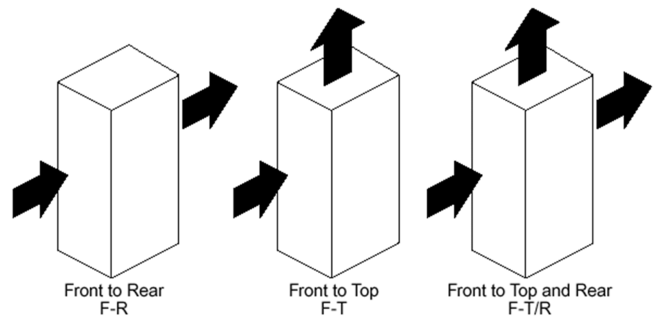

Figure 3 *Recommended equipment airflow protocol.*

Finally, psychometric charts for all environmental classes including NEBS are provided in an Appendix of the guide. These are provided in both SI and IP units to aid the user of these charts. Both recommended (where appropriate) and allowable envelopes are provided for all classes.

Layout

In order for seamless integration between the server and the data center to occur, certain protocols need to be developed, especially in the area of airflow. This section provides airflow guidelines for both the IT/facility managers and the equipment manufacturers to design systems that are compatible and minimize inefficiencies. To ensure this, the section covers the following items:

1. Airflow within the cabinet
2. Airflow in the facility
3. Minimum aisle pitch

In order for data center managers to be able to design their equipment layouts, it is imperative that airflow in the cabinet be known. Currently, manufacturers design their equipment exhaust and inlets wherever it is convenient from an architectural standpoint. As a result, there have been many cases where the inlet of one server is directly next to the exhaust of adjacent equipment, resulting in the ingestion of hot air. This has direct consequences for the reliability of that machine. This guide attempts to steer manufacturers toward a common airflow scheme to prevent this hot air ingestion by specifying regions for inlets and exhausts. The guide recommends one of the three airflow configurations: front-to-rear, front-to-top, and front-to-top-and-rear as shown in Figure 3.

Once manufacturers start implementing the equipment protocol, it will become easier for facility managers to optimize their layouts to provide maximum possible density by following the hot-aisle/cold-aisle concept. In other words, the front face of all equipment is always facing the cold aisle. Figure 4 shows how the inlets would line up.

Finally, the guide addresses minimum practical aisle pitch for a computer room layout. Figure 5 shows a minimum 7 tile pitch where the tile could either be 24 inches or 600 mm.

Table 1 shows how the space is allocated across the seven tiles for either U.S. (24 in.) or Global (600 mm) tiles.

Table 1. Aisle Pitch Allocation

	Tile Size	Aisle Pitch (cold aisle to cold aisle)[1]	Nominal Cold Aisle Size[2]	Max. Space Allocated for Equipment with No Overhang[3]	Hot Aisle Size
U.S.	2 ft (610 mm)	14 ft (4267 mm)	4 ft (1220 mm)	42 in. (1067 mm)	3 ft (914 mm)
Global	600 mm (23.6 in.)	4200 mm (13.78 ft)	1200 mm (3.94 ft)	1043 mm (41 in.)	914 mm (3 ft)

1 If considering a pitch other than seven floor tiles, it is advised to increase or decrease the pitch in whole tile increments. Any overhang into the cold aisle should take into account the specific design of the front of the rack and how it affects access to the tile and flow through the tile.
2 Nominal dimension assumes no overhang, less if front door overhang exists.
3 Typically a one-meter rack is 1070 mm deep with the door and would overhang the front tile 3 mm for a U.S. configuration and 27 mm for global configuration.

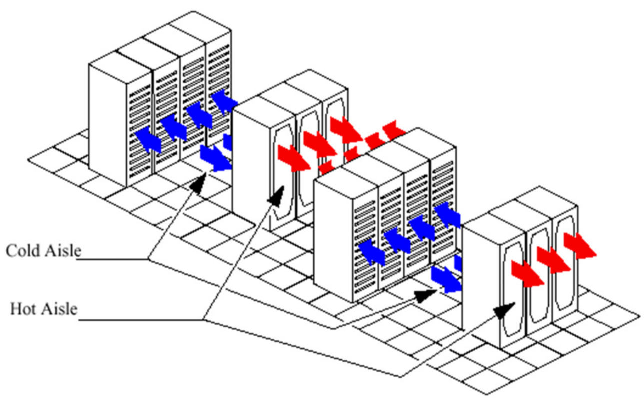

Figure 4 *Top view of a hot-aisle/cold-aisle configuration.*

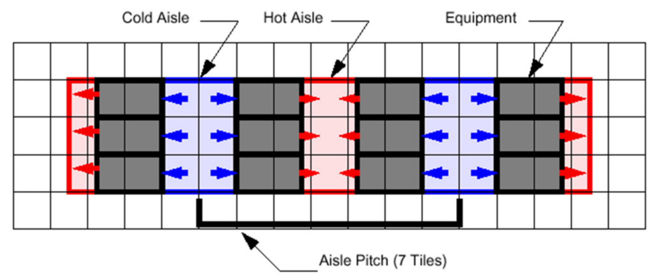

Figure 5 *Minimum hot-aisle/cold-aisle configuration.*

By following these guidelines, equipment manufacturers enable their customers to use the hot-aisle/cold-aisle protocol, which allows them to maximize the utilization of their data centers. In addition, as manufactures adopt the flow directions specified in the guidelines, the swapping out of obsolete servers becomes much less problematic due to the uniform cooling direction.

It is important to note that even if the guideline is followed, it does not guarantee adequate cooling. Although it will provide the best opportunity for success, it is still up to the facility manager to do the appropriate analysis to ensure cooling goals are met.

Power Methodology and Reporting

The ASHRAE guide's heat and airflow reporting section defines what information is to be reported by the information technology equipment manufacturer to assist the data center planner in the thermal management of the data center. The equipment heat release value is the key parameter that will be discussed in this section. Several other pieces of information are required if the heat release values are to be meaningful. These are included in the guide's reporting section and are discussed briefly here.

Currently, heat release values are not uniformly reported by equipment manufacturers and, as a result, site planners sometimes estimate equipment heat loads by using electrical information. Electrical information is always available because IEC 60950 (IEC 1999) and its USA and Canadian equivalent (CSA International 2000) require the maximum power draw to be reported for safety purposes. The safety standard requires rated voltage and current values to be placed on the equipment name-plate label. Electrical power and heat release are equivalent quantities for a unity power factor and are expressed in the same units (watts or Btu/h), but the name-plate electrical information is not appropriate for heat release estimation for data center thermal management purposes.

The design guide states, "Name-plate ratings should at no time be used as a measure of equipment heat release." The first reason is that the name-plate rating is only indicative of a worst-case maximum power draw. This maximum rating will often not be representative of the actual power draw for the equipment configuration to be installed. Second, there is no standard method for defining the maximum power draw. Equipment manufacturers are sometimes motivated to state high rating values so that safety certification current measurements at a rated voltage are well below the rated current. (The safety standard allows the measured current to exceed the rated value by 10%, but this is a situation that manufacturers naturally want to avoid.). The manufacturer may overstate or buffer the rating value to allow the use of higher power components in the future. If the data center planner starts with an inflated nameplate rating and then applies a factor to account for future power increases, the future increase has been counted twice. Third, multiplying a corresponding rated voltage and current value results in a VA or kVA value that must be multiplied by a power factor, which may not be known, to get power in watts or Btu/h. While the power factor is a small adjustment for some modern equipment, not applying the power factor for other equipment may result in another cause of conservative heat load estimates.

To avoid the above problems, the ASHRAE guide defines how heat release information should be determined and reported for data center thermal management purposes. The conditions used to determine the heat release values are specified. They apply to all aspects of the process—measurement conditions, model conditions, and reporting conditions. The conditions are:

- Steady state

  Values based on peak currents are useful for power system design but not for data center thermal management.
- All components in the active state, under significant stress

  The intent is to avoid both unrealistically low values, such as an idle condition, and unrealistically high values. Significant stress in most power or safety test labs is more representative of normal customer operation, while the workload applied in a performance lab may represent a higher than normal workload. Words such as "worst-case" activity were specifically avoided. If heat release values were based on the unlikely condition of all components being in a worst-case activity state, the reported values would be excessively high and the resulting situation would be similar to using name-plate rating values.
- Nominal voltage input
- Nominal ambient temperature from 20°C to 25°C (68°F to 77°F)

  This temperature range is the recommended operating temperature range for Class 1 and Class 2 in Table 2.1 of the ASHRAE document. At higher temperatures, air-moving devices may speed up and draw more power.

  Information technology equipment generally has multiple configurations. Heat release values are to be reported for configurations that span the range from minimum to maximum heat release values. It is acceptable that a heat release value be measured for every reported configuration. It is also acceptable that a predictive model be developed and validated to provide heat release values for multiple configurations. The model would allow the manufacturer to report heat release values for more configurations than could be practically measured. During equipment development, there may be a period when no heat release measurements are available. During this time the model would be based solely on predictions. The ASHRAE document states, "measured values must be factored into the model by the time the product is announced." Appropriate values are measured by the equipment manufacturers as part of the safety certification process, which requires the manufacturer to make electrical measurements. Heat release model validation involves comparing values predicted by the model with measured heat release values for the same configurations.

  The number of tested configurations is not specified, but the required accuracy is defined: the predicted values must be within 10% of the measured values or the model must be adjusted.

Besides heat release values, equipment manufacturers must report additional information for each configuration:

- Description of configuration
- Dimensions of configuration

  Dividing the heat release value by the equipment footprint allows the data center planner to calculate the heat load density in W/m^2 or W/ft^2.
- Weight for configuration

  This is not directly used for thermal management. However, the weight might result in the equipment being spaced apart to meet floor loading requirements. This would result in a decreased heat density, which is important to know for data center thermal management.
- Minimum/maximum airflow characteristics of each configuration in cubic feet per minute (cfm) and cubic meters per hour (m^3/h)

  Unlike the heat release values, which are based on a nominal ambient temperature range, some systems may exhibit variable flow rates due to fan control, which can be dependent upon ambient temperature. For each loading condition, flow rate is to be reported along with the ambient temperature relative to that flow rate. The ambient temperature range should be reflective of the temperature that produces the minimum flow rate as well as the ambient temperature that produces the maximum flow rate. Presumably these temperatures would reflect the allowable ambient extents for which the hardware is designed. The airflow is also reported with all air-moving devices operating normally. For example, if a fan is only powered when another unit fails, the auxiliary unit should be off when determining the airflow value to be reported.
- Airflow diagram

  This can be a simple outline of the equipment showing where the airflow enters and leaves the unit. In the future it may be necessary to provide more information. For example, each inlet and exit airflow arrow may need to be associated with a volumetric airflow value, and the exit airflow arrows may also require a number indicating how much heat the airflow picked up while in the equipment. The goal would be to represent the equipment as a compact model in a data center thermal model.

An example report is included in the guide. It conveys many important aspects of the information to be reported, but it may not be complete for a given product. The example report provides information for minimum, full, and typical configurations. The words "maximum" and "average" configuration were specifically avoided; average particularly may be defined several different ways. It is hoped that the typical configuration used for thermal management purposes will be the same typical configuration used for acoustic measurements. The definition of the typical configuration is left to the manufacturer.

Comparisons to Telecom Specifications

The typical telecommunications central office and the typical data center historically look and operate quite differently. On the electronic equipment side, however, the trend is showing signs of convergence. To meet the demand for new and improved voice, data, and video services, telecommunications service providers are installing a vast amount of non-traditional equipment, which often looks and functions like data-center systems.

Major differences between the telecommunications and data center environments include the following typical central office characteristics.

- Overhead air distribution
- Overhead cable distribution
- Equipment mounted directly on concrete slab
- Tall and narrow equipment frames
- No active humidity control

Despite these significant differences, the data center and telecommunications environments have many thermal challenges and solutions in common. At the core of thermal management is the concept of alternating hot and cold equipment aisles. To maintain the separation of hot and cold aisles, the electronic equipment needs to move the air from the cold aisles to the hot aisles. Airflow from the front-bottom of the equipment to the rear-top works especially well.

Although direct or indirect liquid cooling may be introduced in the future, the current design of equipment is geared toward maintaining air as the cooling medium. Much development and many resources go into making the electronics more compact. One may question the rationale and drive behind this evolution since many telecommunications rooms cannot accommodate the resulting heat output.

ANSI and ETSI documents provide standards for telecommunications central offices in North America and Europe, respectively. The most used de facto standards in North America are Telcordia NEBS GR-63-CORE *Network Equipment—Building Systems* and GR-3028-CORE *Thermal Management in Telecommunications Central Offices* (Telcordia Technologies, Inc. 2002, 2001). The former provides requirements for the physical equipment environment in general, whereas the latter provides specific objectives and requirements for thermal management.

POSSIBLE FUTURE ACTIVITIES

This ASHRAE document reflects the initial recommendation of the ASHRAE TC9.9 participants and is intended to become a live document that would be regularly updated to reflect new data center designs and cooling technologies. Although the power dissipation density trends appear in Figure 2, some new trends such as blade computing indicate that the trends may be steeper and would need to be updated by the ASHRAE TC9.9 team. Furthermore, the first release of the document does not reflect any guidance that may be necessary to integrate the new high-density cooling technologies that are being developed to support the geometric increase in industry standard rack power. Currently, the state of cooling capability of data centers is in the range of 50 to 70 W/ft^2. New installations being designed and constructed today would increase the cooling capacity to approximately 135 W/ft^2. This would allow approximately an average cooling capacity of 3 to 4 KW per rack. The industry is developing products today that when fully populating a rack would dissipate 15 to 25 KW per rack. Thus, there is a need to address the coming product power densities by more aggressive data center cooling design or by other localized liquid and air cooling techniques. In order to enable these technologies to evolve, the future revisions of this document must address the common guidelines for data center and cooling hardware design. This is a significant challenge that would require cooperation between the computer hardware developers and producers, the data center developers and operators, the data center cooling equipment manufacturers, and the rack liquid, air, and refrigerant system developers.

CONCLUSIONS

Although the effort discussed here is significant and extremely useful for the industry, we must recognize that this is only the beginning. It will be necessary for the industry to jointly develop roadmaps into the future that provide a holistic solution for all of the players in the data processing arena.

REFERENCES

ASHRAE. 2003a. *Thermal Guidelines for Data Processing Environments*. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

ASHRAE. 2003b. *2003 ASHRAE Handbook—HVAC Applications*, Chapter 17, "Data Processing and Electronic Office Areas." Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

ASHRAE. 2001. *Standard 127-2001, Method of Testing for Rating Computer and Data Processing Room Unitary Air-Conditioners*. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

CSA. 2000. *Safety of Information Technology Equipment*, CAN/CSA-C22.2 No. 60950-00 3rd ed. CSA International/UL 60950 3rd ed., Underwriters Lab Inc. ANSI/UL 60950-2000 American National Standard.

IEC. 1999. *Safety of Information Technology Equipment*, IEC 60950 3rd ed. International Electrotechnical Commission.

Telcordia Technologies. 2002. GR-63-CORE Network Equipment—Building Systems (NEBS) Requirements: Physical Protection. *Telcordia Technologies Generic*

Requirements Issue 2, April. Piscataway, N.J.: Telcordia Technologies, Inc.

Telcordia Technologies. 2001. GR-3028-CORE Thermal Management in Telecommunications Central Offices. *Telcordia Technologies Generic Requirements* Issue 1, December. Piscataway, N.J.: Telcordia Technologies, Inc.

Uptime Institute. 2000. Heat Density Trends in Data Processing, Computer Systems and Telecommunication Equipment. www.uptimeinstitute.org.

DISCUSSION

Jay Madden, Principal, EYP Mission Critical Facilities, Los Angeles, Calif.: Addressing Table 2.1 in the paper, what do chip manufacturers consider more critical during localized failure: the temperature surrounding the chips or the rate heat rise?

Roger Schmidt: I think both are important, but probably the rate of change may be more critical in the future given the ship technologies that will be developed.

David Copeland, Senior Thermal Engineer, Fujitsu Laboratories of America, Sunnyvale, Calif.: Power density trends show rapid increase from 1997-2005, then much weaker rate of increase. Trends predicted seem to be fairly correct for 1998-2003 product releases. Do we still support predicted trends after 2005?

Schmidt: The ASHRAE TC9.9 committee is in the process of reassessing this trend, and will probably put out an update to this trend chart sometime this year. Basically the server trends seem to be increasing at a much more rapid rate compared to the current trend line, and the storage trends seem to not be increasing at quite as fast a rate as what the trend line shows.

Rebecca Perry, Engineer, Sun Microsystems, San Diego, Calif.: Is it possible to design multivendor products with consistent building-block-like cooling characteristics (not just airflow direction) to reduce complexity to the end user and make product integration in a data center easier and more reliable?

Schmidt: I think it is, but it will take an industry group to drive this commonality. I think the ASHRAE TC9.9 committee can play an important role in this commonality.