ASHRAE TC9.9

Data Center Networking Equipment – Issues and Best Practices

Whitepaper prepared by ASHRAE Technical Committee (TC) 9.9
Mission Critical Facilities, Data Centers, Technology Spaces, and Electronic Equipment

Table of Contents

Executive Summary

New thermal design and installation practices are needed to prevent the possibility of over-heating and loss of functionality of networking equipment in a data center environment. Commonly accepted design and installation methods, such as placing a top of rack switch behind a blanking panel, can cause networking equipment to over-heat, lose functionality, and even compromise data integrity. New energy saving trends, such as higher data center air temperatures and economization, are further taxing the thermal margins of networking equipment. A combination of design and installation recommendations are proposed that will allow for seamless and reliable integration of networking equipment in a data center environment with little or no need for thermal engineering on the part of the user and no concerns of over-heating or compromise of functionality.

ASHRAE TC9.9 recommends new networking equipment designs draw cooling air from the front face of the rack with the air flow direction from the front of the rack to the rear and the hot exhaust exiting the chassis at the rear face of the rack. This front to rear cooled equipment should be rated to a minimum of ASHRAE Class A3 (40°C) and preferably ASHRAE Class A4 (45°C). The development of new products which do not adhere to a front to rear cooling design is not recommended. It is recommended that networking equipment, where the chassis doesn't span the full depth of the rack, have an air flow duct that extends all of the way to the front face of the rack. ASHRAE TC9.9 recommends the equipment be designed to withstand a higher inlet air temperature than the data center cooling supply air if: a) the equipment is installed in an enclosed space that doesn't have direct access to the data center air cooling stream, or b) the equipment has a side to side air flow configuration inside an enclosed cabinet.

Networking equipment manufacturers should provide very specific information on what types of installations their equipment is designed for. And users should follow the manufacturer installation recommendations carefully. Any accessories needed for installation, such as ducting, should either be provided with the equipment or should be readily available. By following these recommendations the risk of equipment over-heating can largely be avoided and the compatibility of networking equipment with other types of equipment in rack and data center level solutions will be significantly improved.

1 Introduction

This paper is written for a broad audience that includes data center power and cooling experts as well as IT and networking specialists. Sections of the paper may be basic for some audience members. The reason for this is to provide a wide base of background information that bridges any gaps in understanding and provides a common framework for understanding the proposed networking thermal issues and best practices.

Significant changes are taking place in data centers that will affect how networking equipment is designed and deployed both now and in the future. For example, many data center applications require networking equipment to be deployed as part of a rack level solution. Rack level solutions can create an interaction between the networking equipment and the thermal behavior of other types of IT equipment in the rack. The exhaust temperature of current generation servers has risen significantly as fan speeds are being reduced to save energy. In many common data center rack level installations, the networking equipment takes its cooling air flow from the rear of the rack where the air temperature is largely determined by the exhaust temperature of the other IT equipment. Data center operating temperatures and cooling technologies are also changing. For more information on new and emerging data center cooling technologies such as air-side economization, water-side economization, liquid cooling and the efficient use of air conditioning, please consult the books in the ASHRAE Datacom series [1-4]. Traditional operating ranges of 15-20°C are giving way to warmer temperatures, some as high as a recommended temperature as high of 27°C. The adoption of economization (both air and water-side) is growing. In a heavily economized data center, the air inlet temperature of the IT equipment will be determined by the temperature of the outdoor air and can vary widely depending on the time of day and season of the year. Even in a conventional HVAC controlled data center that is operated within the ASHRAE recommended range of 18 – 27°C, it is possible to have an installation where the networking equipment exceeds its data sheet temperature rating under normal operating conditions. Exceeding the maximum rated temperature is not allowed and it may impact data integrity or even cause a loss of functionality. Recently, a major data center operator [5] made a public statement that they believe the thermal capability of the networking equipment was a weak link in their fresh air cooled data centers. Given this backdrop of changes to data center environments and IT equipment thermal design, this is an opportune time to revisit the thermal design and installation best practices for data center networking equipment.

The purpose of this paper is to review the current state of data center networking equipment thermal design and installation practices. Ideal and non-ideal designs and installations will be compared to highlight common thermal and air flow issues. The paper will conclude with a set of recommended thermal design and installation best practices for networking equipment in a data center environment.

1.1 Market Segments for Networking Equipment

The networking equipment market can be broken out by the customer segments it serves and their respective usage lifetime expectation as shown in Table 1. Data center networking equipment primarily comes from the Enterprise category. However, quite a bit of Service Provider (SP) equipment and even some Small Office Home Office (SOHO), branch, medium size office equipment can also be found in data centers. Interestingly, much of the networking equipment today is being designed and sold so it can be deployed across multiple market segments.

Table 1 Market segments for networking equipment

Market Segment	Description	Product Example	Expected Service Lifetime (years)
Consumer	Customer premise equipment (CPE) used with service provider services. They include telephone, Digital Subscriber Line (DSL) or cable modems, set-top boxes, etc.	Home switches and routers, set top boxes	3-5 years
SOHO, branch, medium office	Typically 1-10 employees for Small Office Home Office (SOHO), and for branch and medium office, typically the equipment is housed in a designated room in the office	Integrated services routers, entry-level enterprise-class fast Ethernet switches	5-7 years
Enterprise	For large corporations that house IT equipment in data centers	High performance modular switches for unified communications	7 years
Service Provider (SP)	Telecom companies that house switching equipment in central offices	Telecom switching equipment	>10 years

1.2 Basic Networking Hardware Functionality

A computer network, or simply a network, is a collection of computers and other hardware interconnected by communication channels that allow sharing of resources and information. Networking equipment is equipment that facilitates the interconnection of devices and the sharing of data both within the data center and beyond. A switch is a device that receives a message from any device connected to it and then transmits the message only to the device for which the message was meant. One of the simplest and most common types of networking hardware is a top of rack (TOR) switch. A photo of a top of rack switch with the cover removed and the major functional blocks labeled is shown below in Figure 1.

Fig. 1 Top of rack switch with cover removed.

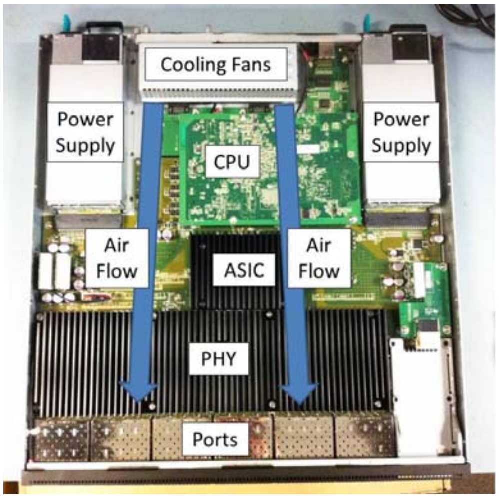

Data enters through the ports and travels to the PHY or physical layer chips. The Application Specific Integrated Circuit (ASIC) chip decides where that data needs to go and sends it back out through the correct port. The Central Processing Unit (CPU) controls both the PHY and the ASIC. The CPU can take data off of the network, process it, and then send it back out onto the network. The CPU is considered out of band because it is off the network and does not function at the full I/O speed. Encryption is an example of data processing done by the CPU. One of the key functions of the ASIC is to handle Priority Flow Control (PFC) rules. PFC rules prioritize which data is sent first and which data can wait. One example of where PFC rules come into play is the transmission of video. Video is nearly always given a high PFC priority because any delays in transmission could cause a glitch that would be noticeable to the end user. The ASIC also performs valuable functions such as blocking network traffic and stopping a flood of data, if this should occur.

The functional elements of large data center or telecom central office networking equipment are very similar to those of the small top of rack switch except that the size and scale of the system is larger. Also, larger switches tend to use optical ports instead of copper Ethernet cable type ports.

1.3 Abstraction Layers Used in Networking

Functions of a networking system are characterized in terms of abstraction layers. An abstraction layer (or abstraction level) is used as a way of hiding the implementation details of a particular set of functionality. A majority of data center networking equipment and software uses the ISO/IEC 7498-1 standard Open Systems Interconnection (OSI) 7-layer model [6] for computer network protocols. In this model, similar communication functions are grouped into logical layers. A layer serves the layer numerically above it and is served by the layer numerically below it. A description of the purpose and functionality of the more commonly used OSI layers is given below:

- **Layer 1** = is generally considered the physical layer. A network hub, or repeater, is an example of a simple Layer 1 network device. Any packet entering a port of the hub will be "repeated" on every other port, except for the port of entry.

- **Layer 2** = is generally used for basic point A to point B routing and is considered the data link layer. A network bridge is a good example of a Layer 2 device. The bridge may interconnect a small number of devices in a home or office. Single bridges can also provide extremely high performance for specialized applications such as storage area networks.
- **Layer 3** = is intended to handle higher level functions such as priority flow control and spanning tree protocol.
- **Layer 4** = the exact meaning of layer 4 is vendor-dependent. However, it almost always starts with a capability for network address translation.
- **Layers 5 & 6 –** these layers are not commonly used.
- **Layer 7** = includes the graphical user interface layer of your personal computer or server. Layer 7 switches may distribute loads based on Uniform Resource Locator (URL) or by an installation specific technique to recognized application-level transactions. A layer 7 command will typically span lots of commands across many different layers.

1.4 Common Types of Networking Equipment

Common types of networking equipment, their functions and the OSI layers they use are given in Table 2 below.

Table 2 Common types of networking equipment and their functions.

Equipment Type	OSI Layers	Description
Repeater	OSI layer 1	Amplifies or regenerates the digital signals received while sending them from one part of a network into another.
Hub	OSI layer 1	Connects multiple Ethernet segments, making them act as a single segment.
Bridge	OSI layer 2	A device that connects multiple network segments along the data link layer.
Switch	OSI layer 2	A device that allocates traffic from one network segment to certain lines which connect the segment to another network segment. Unlike a hub, a switch splits the network traffic and sends it to different destinations rather than to all systems on the network.
Router	OSI layer 3	A specialized network device that determines the next network point to which it can forward a data packet towards the final destination of the packet
Gateway	OSI layers 4 and 7	Device is placed at a network node and interfaces with another network that uses different protocols.

1.5 Typical Data Center Network Topology: An Overview

Design and deployment of a network in a data center is not a trivial task. However there are a few implementation concepts that are common in many data center network deployments which are important to understand when considering the environment in which the networking equipment functions. For general information on networking technology please consult references [7-10].

First of all, networks tend to be designed in a hierarchical manner, with many devices (in the case of a data center, servers and storage devices) connected to a switch which is connected up to another switch at the next level of the hierarchy, and so on. Another common topology is a mesh configuration, in which many peer network switches are connected to one another to form a single level of the hierarchy. For the purposes of this overview, we will consider three different levels of a network hierarchy: core, distribution, and edge.

The core network function can be thought of as the gateway through which all data entering and exiting the data center must pass. As such, the core network equipment is connected either directly or indirectly to every device in the data center. The core switch is also connected to a service provider (SP), which is the "pipe" through which all data passes from the data center to the internet.

The distribution level of the network acts as an intermediate level between the core and edge levels of the network, and as such can offload some of the work the core network equipment needs to do. Specifically the distribution level is useful for passing data between machines inside the data center, or aggregating ports to reduce the number of physical ports required at the core.

The edge network level consists of switches that connect directly to devices that are generating or consuming the data, and then passing the data up to the distribution or core. A data center network implementation may have all three of these levels, or it may combine or eliminate some of the levels.

Figures 2 through 4 are simplified diagrams showing how these concepts might be implemented in a data center, and where particular types of networking equipment may be located. Understanding where the equipment is located will help the equipment supplier and those responsible for deploying the network understand the environment that the network equipment will be exposed to.

For each of the types of networking equipment shown in the figures, a preferred infrastructure configuration will be suggested. The infrastructure configurations will primarily make reference to the type of rack cabinet to be used and the direction that the ports face in a standard hot aisle/cold aisle rack configuration.

In general, two types of racks cabinets are widely available in the industry. The first is a standard EIA 19 inch rack [11] enclosed in a cabinet which is approximately 24 inches wide. This will be referred to as a *server rack*. The second is a standard 19 inch rack in an enclosed cabinet which is typically several inches wider than the cabinet of a server rack. This will be referred to as a *networking rack*. The advantage of the wider networking rack is twofold: it allows a significant space for routing of large bundles of cables vertically next to the rack columns, and it provides extra plenum space for side to side cooled networking equipment.

Fig. 2 Example of a simple network implementation typical of a small data center.

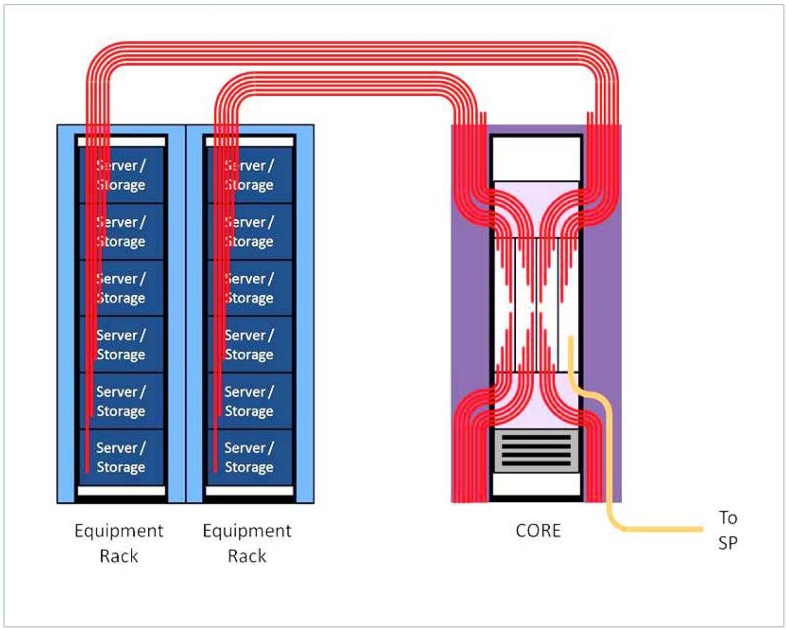

Figure 2 shows a simple network implementation in which the core, distribution, and edge level functions are all performed by the core network equipment, the equipment which is directly connected to the SP. This type of configuration is common for small data centers, but may be utilized for larger data centers when a very flat network hierarchy is desirable. Note that in the case of larger data centers, the load on the networking equipment can be substantial with this configuration, both in terms of port count and data throughput.

Because there is a single connection to the SP, the core networking equipment for the implementation shown in Figure 2 must all be located together in one physical location in the data center. For this type of implementation, the preferred infrastructure configuration is to use networking style racks with wider (vs. traditional EIA standard 19" server style racks) enclosed cabinets for the core networking equipment and server racks for the server and storage equipment. The ports should face the cold aisle for the core networking equipment rack and if side to side cooled networking equipment is used, ducting should be installed in the plenum space afforded by the networking rack to draw air in from the cold aisle and exhaust it to the hot aisle. Though not shown in the figure, it may be advantageous to utilize patch panels to allow for ease of network configuration from the cold aisle side.

Fig. 3 Example of a network implementation typical of a medium to large size data center.

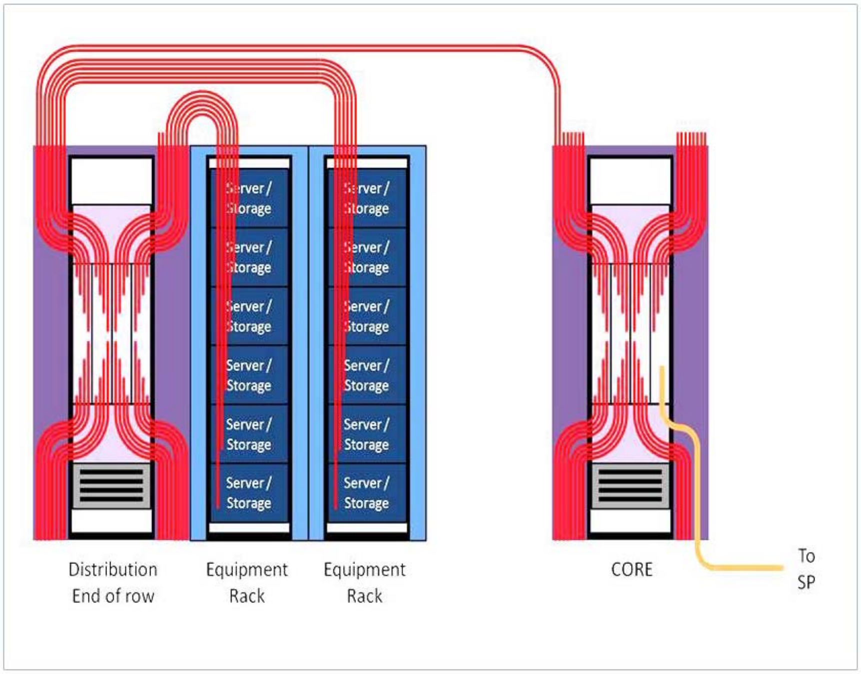

Figure 3 shows a slightly more complicated network implementation in which a distribution level has been added to the network in the form of an end of row network cabinet. The distribution equipment offloads the core equipment in both throughput and port count. In a medium to large size data center with multiple equipment rows it is common for all of the distribution networking equipment to be connected together on the same level of hierarchy to form a mesh (not shown).

The racks above in Figure 3 are drawn to show networking topology and not how they would be physically arranged along the rows of a data center for cooling. The core and distribution switches have been drawn facing the cold aisle whereas the rear of the server racks are shown which would normally face the hot aisle. The preferred infrastructure configuration for the end of row cabinet should be the same as for the core networking equipment. It should utilize a networking rack with the ports facing the cold aisle, ducting for side to side cooled networking equipment, and patch panels for ease of network configuration.

Fig. 4 Example of a typical network implementation with top of rack switches.

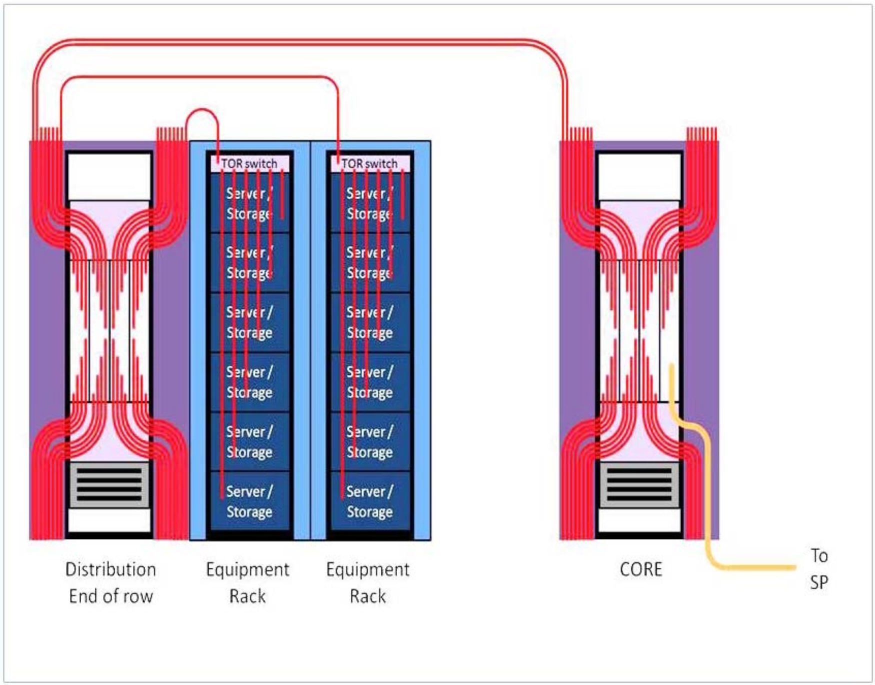

Figure 4 adds another level to the network implementation in the form of edge switches at the top of each equipment rack, commonly referred to as top of rack (TOR) switches. The TOR edge networking equipment offloads the distribution equipment in both throughput and port count, similar to the way the distribution networking equipment offloads the core networking equipment. As with the display in Figure 3, the racks in Figure 4 are drawn to show networking topology and not how they would be physically arranged along the rows of a data center for cooling. The core and distribution switches have been drawn facing the cold aisle whereas the top of rack switches are shown facing the hot aisle. In a typical data center installation, the cable ports of core and distribution switches face the data center cold aisle whereas the ports on small top of rack switches typically face the hot aisle since the server networking ports are typically located at the rear of the servers in the hot aisle.

The preferred infrastructure configuration for the TOR switch is different than for the core and end of row equipment. The TOR switch typically resides at the top of the equipment rack which is usually a server rack. The ports should be facing the rear of the rack (hot aisle) because that is typically where the network connections are located on server and storage devices. So the preferred airflow direction for the TOR switch is front to rear cooling, where the rear of the TOR switch is where the network port connections exit.

2 Survey of Maximum Temperature Ratings

A comprehensive benchmarking survey was carried out to determine typical maximum temperature ratings for data center networking equipment, whether there is any standardization of maximum temperature ratings, and do these maximum temperature ratings align to industry standards and the ratings of other types of IT equipment. Results of the benchmarking survey are shown in Figure 5 with benchmarking survey results for servers and disk storage equipment included for comparison.

Fig. 5 Maximum temperature rating benchmarking survey results.

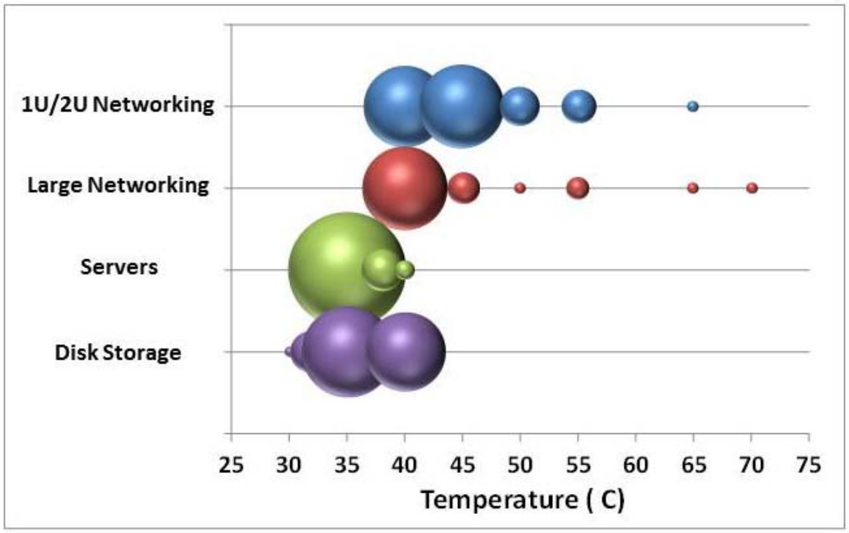

The benchmarking results above represent a survey of 9 of the leading networking equipment suppliers (nearly 400 models of equipment), 6 of the leading server suppliers (nearly 300 models), and 7 of the leading disk storage suppliers (nearly 200 models).

For a given horizontal category, the sum of the area of all of the bubbles represents 100% of the product models and the relative area of individual bubbles scales with the percent of product models rated at that temperature. Per the data in Figure 5, there doesn't seem to be a standard or even a consensus for a single maximum temperature rating for small (1U/2U form factor typical of top of rack switches) networking equipment. The majority of the equipment is rated at either 40 or 45°C but there are some models rated as high as 50 and 55°C. In contrast, large networking equipment, with a few exceptions, has a common design point of 40°C. The 40°C design point may be driven by the requirements of NEBS/Telcordia GR-63-CORE [12] as most models of large networking equipment are typically NEBS certified. Servers are fairly homogeneous with 87% of server models rated at 35°C with 11% rated at 38°C and 2% rated at 40°C. Disk storage arrays don't have a single common maximum temperature design point – with a few exceptions, they are almost evenly divided between 35 and 40°C. In general, networking equipment is rated 5 – 10°C higher than servers.

3 Cooling Design of Networking Equipment

3.1 Common Air Flow & Mechanical Design Configurations

From an air flow and mechanical design standpoint, there are several common physical design configurations for data center networking equipment. However, before embarking on a discussion of air flow it is important to establish standard definitions for the faces of the chassis of networking equipment and the rack it is mounted in such that we can use to reference air flow direction. The 2008 ASHRAE Thermal Guidelines for Data Processing Environments [13] defined a standardized nomenclature for defining the air flow paths for server equipment as shown below in Fig. 6. Table 3 further defines cooling air flow directions for networking equipment. NEBS/Telcordia GR-3028 [14] contains definitions for cooling air flow directions but for equipment that is used in a telecom environment. The ASHRAE definitions will be used throughout this document.

Fig. 6 ASHRAE standard nomenclature for air flow paths through server equipment [13]

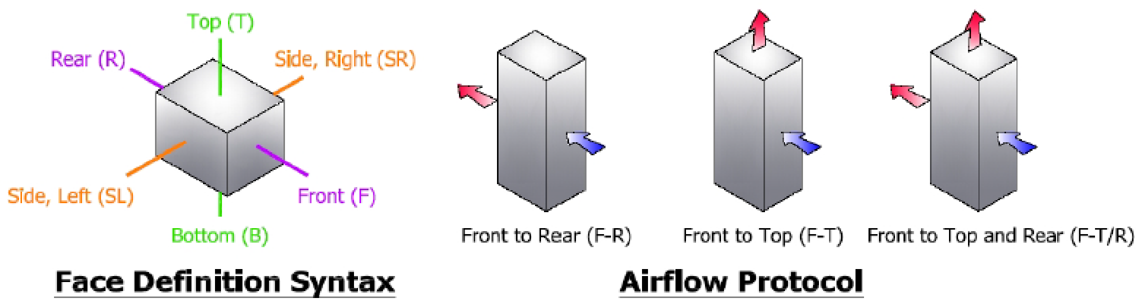

Table 3 Definition of cooling air flow directions

Face of Chassis	Definition
Front Side (F)	Side of the networking equipment chassis facing the front of the rack. The front of the rack is typically adjacent to the cold aisle or cold air supply.
Rear (R)	Side of the networking equipment the faces the rear of the rack. The rear of the rack is usually where the hot exhaust of most IT equipment, such as servers and storage equipment, is directed. The rear of the rack is adjacent to the hot air exhaust.
Cable Port Side	Side of the equipment where the LAN cables plug into the cable ports. In small 1U or 2U data center network switches, this side typically faces the rear of the IT equipment rack.
Side Right (SR)	The right side of the chassis is defined from the point of view of one standing in the cold aisle or facing the front of the IT equipment rack.
Side Left (SL)	The left side of the chassis is defined from the point of view of one standing in the cold aisle or facing the front of the IT equipment rack

Several of the most common physical design configurations for data center networking equipment are shown below in Figure 7.

Fig. 7A Large switch, typically full and half size rack, with front to rear air flow.

Side of Rack View
(side panel of rack & chassis removed)

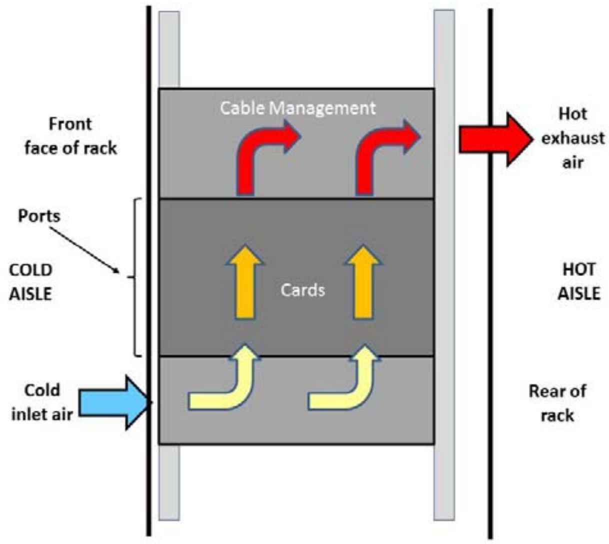

Front of Rack View

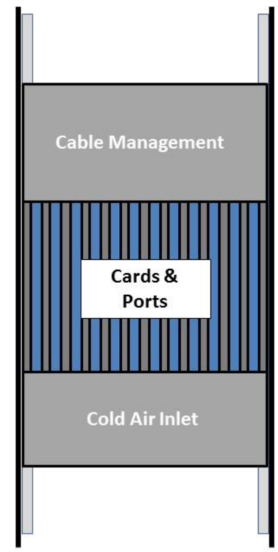

Note that the air flow is actually "S" shaped - cool air enters horizontally at the bottom front of the chassis and then travels vertically along the cards before travelling horizontally as it exits towards the rear at the top of the chassis. The network ports come off of the front facing edge of the cards and are located in the front middle of the system.

Fig. 7B Mid size switch with side to side air flow.

Side of Rack View
(side panel of rack and chassis removed)

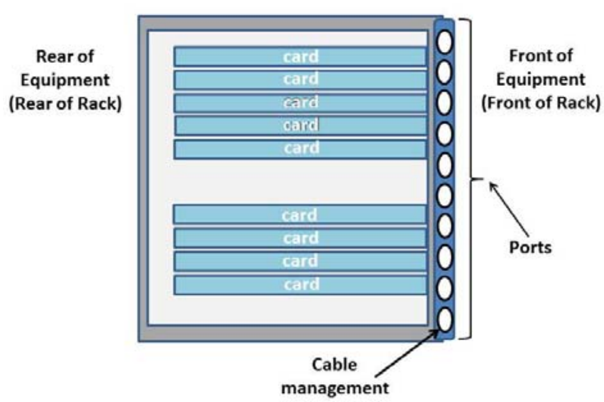

Front of Rack View

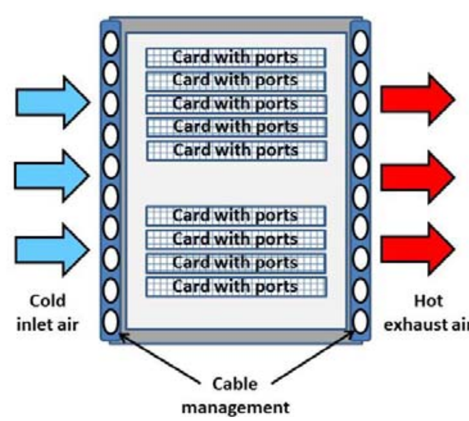

Note how the cards are oriented horizontally and the side to side cooling air flow stream travels parallel to the cards.

Fig. 7C Small networking equipment with “S” shaped air flow.

Top View
(top panel of chassis removed)

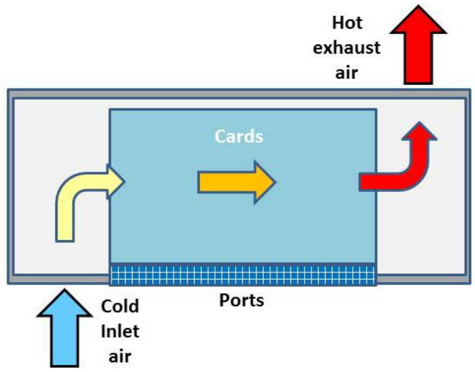

Port Side View
(typically faces front of rack)

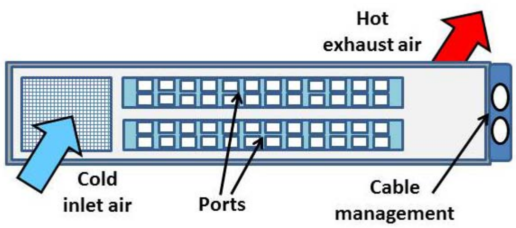

Cool air enters through the left front side of the chassis and then travels horizontally along the cards until it is directed out the right rear side.

Fig. 7D Small 1 or 2U switch with side to side air flow.

Top View
(top panel of chassis removed)

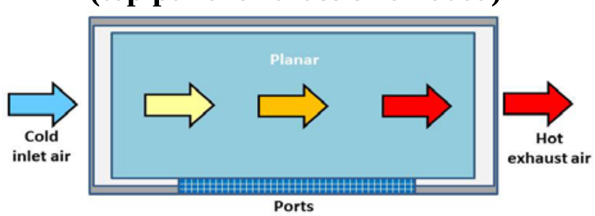

Port Side View
(typically faces rear of rack)

Fig. 7E Small 1 or 2U switch with front to rear air flow.

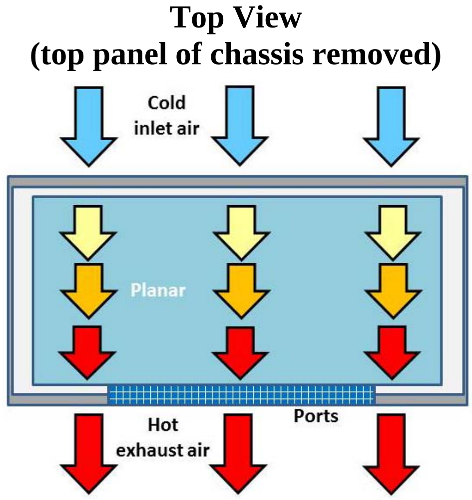

Note: some of these switches also have a reverse air flow feature.

Fig. 7F Small form factor equipment with only natural convection cooling.

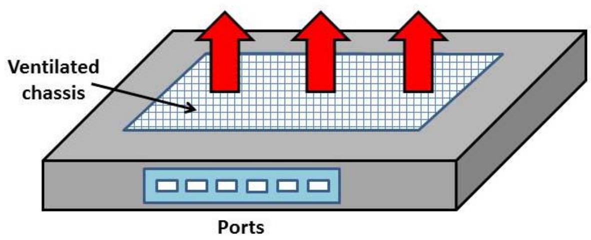

The designs shown above in Figure 7 are not all inclusive – they are only intended to show some of the most common designs.

Some of the common elements these designs share are:

- Air flow is usually routed parallel to the cards
- Cable management is designed so as not to impede air flow
- Network ports are typically located for user convenience, not for optimal thermal design or air flow

3.2 Origin of Networking Equipment Thermal Designs

Typically, server and storage equipment designed for the data center utilizes front to rear airflow for system cooling. This airflow paradigm works well with the hot aisle/cold aisle arrangement of most data centers. Much of the networking equipment on the market today uses side to side or side to rear airflow. Some of the reasons networking equipment has evolved with thermal designs that are unique from other types of IT equipment are:

1) **Higher port density in same/smaller form factor**. In a 1U form factor system that has a large number of ports (e.g. 48) almost an entire side of the chassis is taken up with ports

leaving little or no room for air flow. Because side to side air flow designs use the sides of the chassis instead of requiring a larger port side, they typically use 30% less space than a comparable front to rear air flow configuration.

2) **Improved air flow**. Side to side cooling using the sides of the chassis, provides the surface area needed for larger inlet and outlet grills that allow more air flow with less back pressure. Historically, many types of racks/cabinets have either been open or have had a significant amount of space along the sides of the IT equipment for air flow.
3) **Cable management.** Networking equipment has significantly more cabling than most other types of IT equipment. A large number of networking cables can restrict cooling air flow and, if the cables form a tight droop over a cooling vent, they can further restrict air flow. Another reason why side to side air flow has been favored is it routes the cooling air flow through the side panels where there is no cabling.

Fig. 8 Photo showing "cable sprawl" from networking equipment.

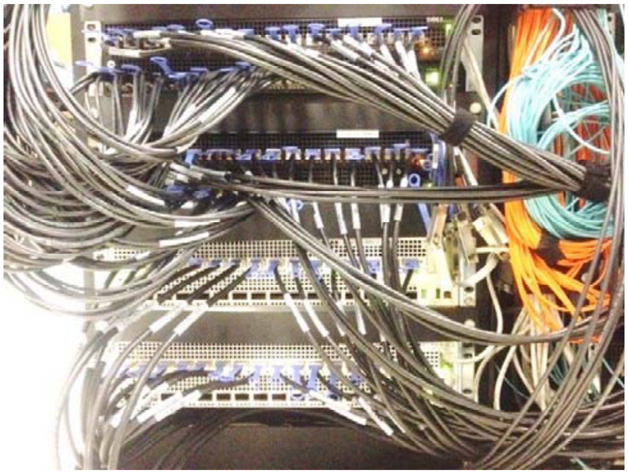

4) **Air filtration requirements & physical access to cards and cabling.** Air filtration is a requirement for most networking equipment for service provider and telecom applications. The air filter has to be designed to be user accessible and have replaceable or cleanable media yet it can't interfere with access to ports and cards or interfere with the visibility of LED indicators on those cards. Thus, in most networking equipment the space used for air flow is separate from the space used for physical access to cards, cables, power switches, etc. These constraints tend to favor side to side air flow or air flow designs other than a front to rear air flow.
5) **Data center evolution.** Data centers have been evolving from a model in which networking equipment was in its own centralized space to one in which it is more distributed throughout the data center. When networking equipment was installed in a dedicated area the unique airflow requirements of networking equipment did not present the challenge that it does when distributed with other IT equipment.

The equipment design priorities above that have led IT networking equipment manufacturers to the thermal designs that are currently available may need to be revisited to integrate seamlessly with the current state of rack level and data center cooling designs and with the possibility of higher inlet ambient temperatures.

3.3 Ideal Designs & Installations

If networking equipment is going to keep pace with changes to data center temperatures and the need to be integrated as part of a rack level solution, the design and installation best practices for networking equipment will need to evolve. Several ideal data center networking equipment design and rack level installations are show below in Figures 9 and 10.

Fig. 9 Small (1U/2U form factor) rack mounted networking equipment showing ideal air flow.

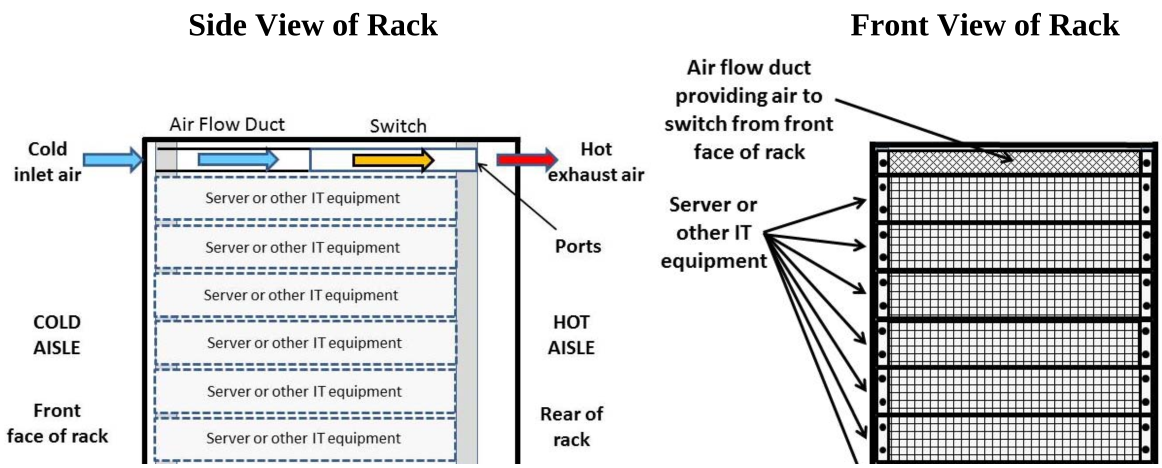

This can be achieved by a combination of a switch that has front to rear air flow and duct work that extends the switch to the front face of the rack.

Fig. 10 Large half rack networking equipment showing ideal air flow.

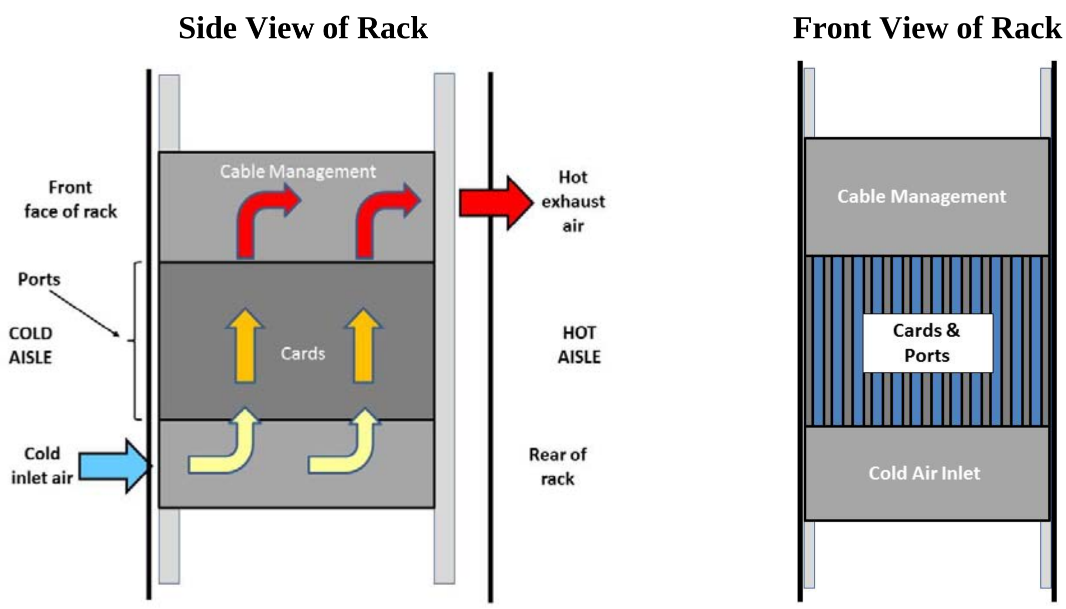

Although Figs. 9 and 10 appear to be different types of installations and different form factors of equipment, the equipment designs and installations above are considered ideal because they share several important points in common:

- Front of rack to rear of rack air flow
- Cooling air flow is taken from the front panel of the rack instead of from inside the rack where mixing of hot and cold air streams can occur
- Hot air is exhausted at the rear face of the rack
- Cold & hot air streams are isolated from one another inside the rack and inside the equipment to prevent mixing
- There are no bypass paths around the networking equipment or through the rack to allow cold air to mix with hot air
- The design of the air flow paths is simple. In the case of the small switch, the air flow path is linear with very little impedance to the air flow. For the large switch, the air flow does have two bends but the equipment is designed to function with the added impedance those bends create.

3.4 Non-Ideal Designs & Installations

Many of the issues with current data center networking equipment design and installations are typified by the illustration in Figure 11 and the photos in Figure 12 below.

Fig. 11 Small 1U form factor switch mounted behind a blanking panel.

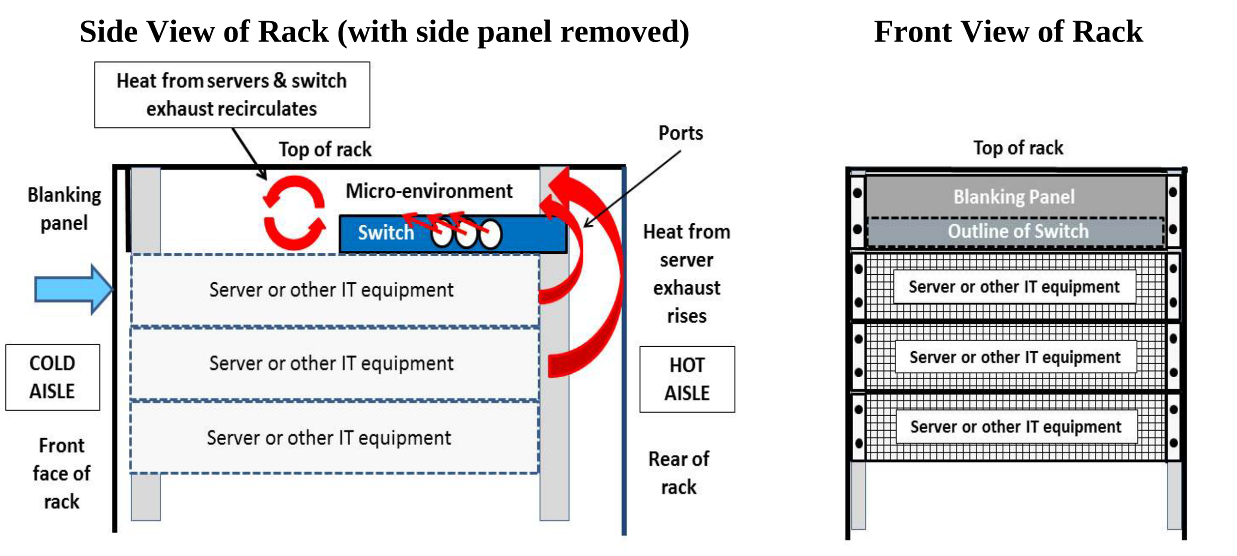

Note how the blanking panel creates a closed micro-environment within a small space that is cut off from the cold aisle where the air temperature is determined by re-circulated hot air from the server and switch exhausts. In this installation configuration, the pre-heated air inside the micro-environment is what the switch must use for its cooling.

Fig. 12 Infrared images of switch installed behind a blanking panel.

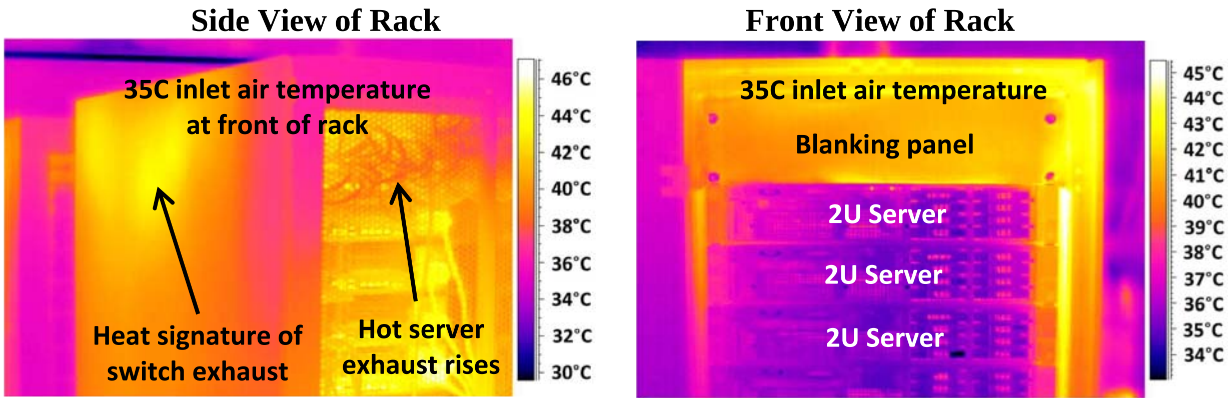

Note how you can see the heat signature of the switch exhaust on the side panel of the rack. Note also how hot the server exhaust is and how it tends to rise to the top of the rack. In the front view of the rack, the blanking panel glows orange indicating how much heat is being trapped behind it.

The blanking panel serves a purpose – it prevents hot rear exhaust from escaping through the front of the cabinet and mixing with cold air where it could compromise the cooling air supply of IT equipment in the rack such as servers. However, the blanking panel creates a micro-environment that is potentially troublesome for the switch product as it blocks access to cold air. This type of rack installation forces the switch to draw its cooling air from the hot re-circulation region inside the rack which is determined by the exhaust temperature of the other IT equipment in the rack. Current generation server exhaust temperatures have risen significantly because the thermal design of these servers slows the fan RPM down in order to save energy. The graph below in Fig. 13 shows examples of typical exhaust temperatures for current generation 1 and 2U rack servers.

Fig. 13 Server exhaust temperature as a function of inlet ambient temperature.

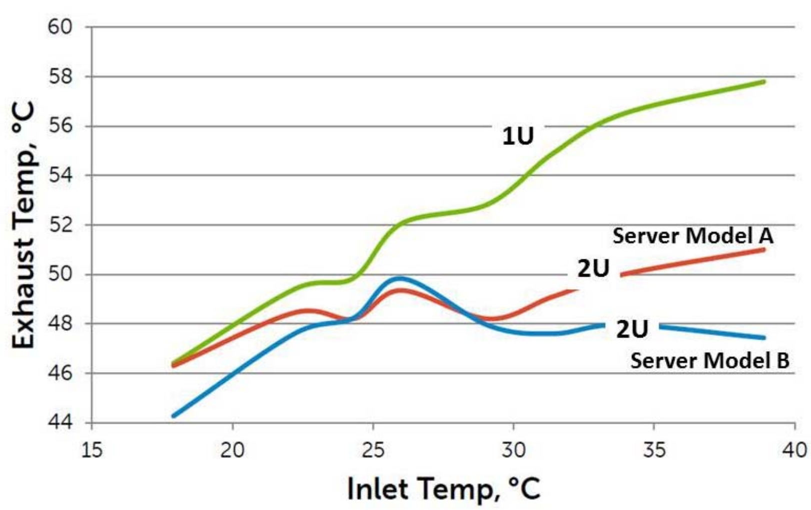

With an inlet temperature of 27°C, which is still within the ASHRAE recommended range, the exhaust temperature can reach 52°C, which is above the maximum temperature rating of many small form factor switches.

What is interesting to note is much of the server exhaust temperature range of (44 – 58°C) is above the 40-45°C rating of a typical network switch. Thus, in the type of networking equipment installation shown in Figure 11, a thermal environment could be created inside the rack/cabinet that exceeds the maximum temperature rating of the switch even under fairly benign front of rack IT inlet temperature conditions.

Another example of a non-ideal installation is shown below in Figure 14 where networking equipment is installed in open 2-post style racks. The side to side air flow design of the equipment directs the hot exhaust of the equipment in the neighboring rack into the equipment cold air inlet. A baffle needs to be installed between each rack of equipment to mitigate the problem. However, even with a baffle, the installation is inefficient because the hot and cold exhaust mix with the room air and there is little or no containment of hot and cold air streams.

Fig. 14 Front view of side to side cooled networking equipment mounted in 2-post open style telecom racks.

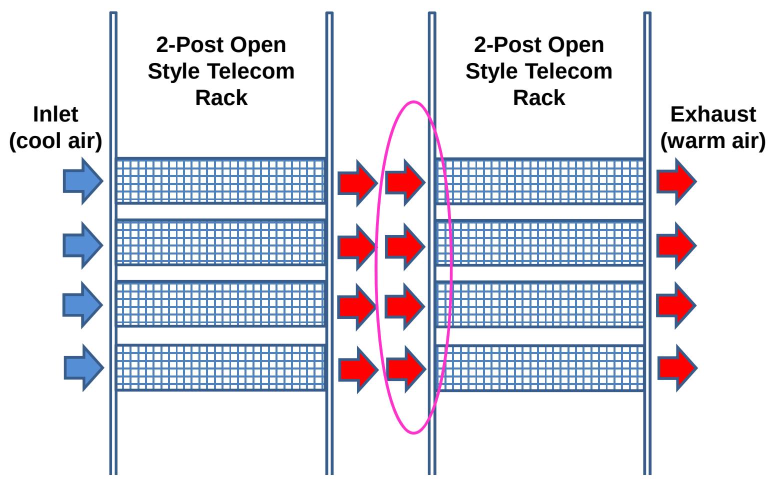

The side to side cooling air flow design of the equipment causes the warm exhaust from one piece of equipment to feed into the inlet for the cool air stream of the equipment in the adjacent rack. This installation will cause the second piece of equipment to operate at much higher inlet temperatures than intended.

Non-ideal design practices to be avoided are:

- **Hot and cold air flow streams to the networking equipment are not isolated from one another and are not contained.** In the case of the switch installed behind a blanking panel, the switch is attached to the rear frame of the rack and the switch chassis occupies only a portion of the depth of the rack. This creates a bypass path for hot and cold air to mix and flow through the rack around the switch.

- **Localized micro-environments are created** – micro-environments created within a small space raise the cooling air inlet temperature of the networking equipment well above the nominal temperature of the data center air. In one case, the micro-environment is created by a blanking panel and in the other case it is created by installing equipment in the heated air stream of adjacent equipment.
- **No direct access to the cold air stream provided by the data center** - the equipment doesn't have direct access to the cool air stream provided by the data center air handling system. Thus, the inlet air temperature to the IT equipment is not directly determined by the data center environmental controls.
- **Side to side air flow configuration is non-standard** – most other types of IT equipment in data centers, such as servers and storage, have a front to rear air flow configuration. Networking equipment that has a side to side configuration is non-standard and can pose problems when integrating different types of IT equipment into a rack level solution. It should be noted that, with careful equipment design and installation, side to side air flow equipment can be made to work.

4 Equipment Power and Exhaust Temperature

Although the rate of power increase of IT equipment generation to generation has moderated since 2005, the power consumption for networking and communications equipment is still increasing as shown below in Table 3. The compute/storage performance/watt continues to increase dramatically across all IT products. Cooling contributes to this improvement with much better fan control algorithms that permit fan speed reductions resulting in improved energy efficiency. When the IT equipment moves less cooling air through the equipment for a given heat load, the temperature of the exhaust air will increase. Where this becomes an important consideration for networking equipment inside an enclosed rack is when the exhaust air from servers or storage equipment is ingested into adjacent networking equipment. And, under some circumstances, the heat generated by the networking equipment itself can contribute to the temperature rise inside the rack further raising the temperature of the air into the networking equipment (see Figs. 15 and 16 below). This problem is compounded by rising data center temperatures, by an expansion of the data center operating ranges, and by wider adoption of air-side and water-side economization.

Table 4 Power trends in watts/ft^2 of non-standard planform equipment.

Type	Range of Average Heat Loads	Range of Footprints (ft^2)	Heat Load Per Product Footprint (Watts per sq.ft.)		
			2010	2015	2020
Storage Servers	+/-15%	6 - 13.5	700	850	1100
Tape Storage	+/-30%	10 - 12	200	200	200
Communications	+/-20%	6 - 12	2000	2550	3000

Networking equipment typically has a much higher power density than storage or servers [15].

Fig. 15 Example of relationship between inlet and outlet air temperature for a 10U edge router.

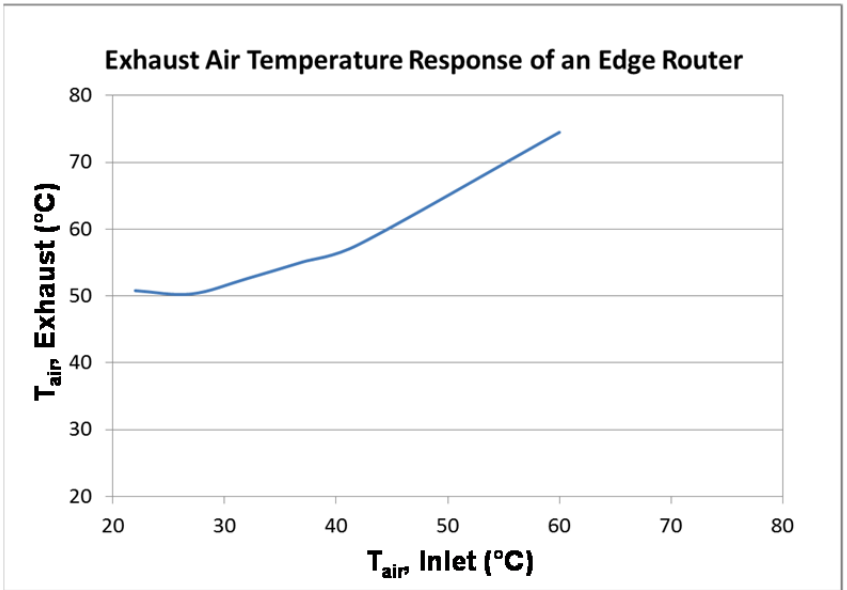

Note at 20°C inlet the air temperature rise is nearly 30°C whereas for an air inlet temperature of 60°C the air temperature rise is closer to 15°C. At 60°C the fans are running at high speed and moving much more air than at 20°C.

Fig. 16 Example of exhaust temperature rise for a small top of rack switch.

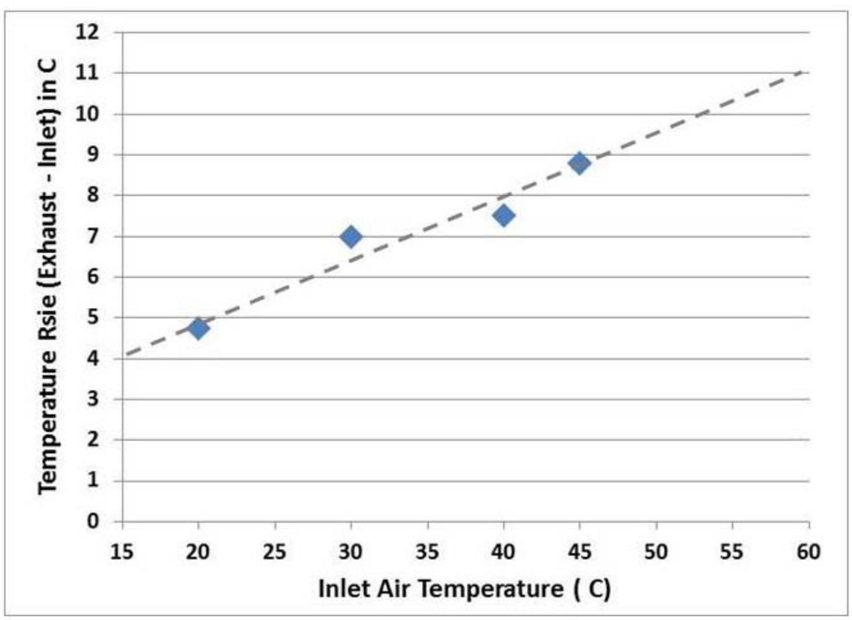

5 Environmental Specifications

The most commonly used environmental thermal guidelines and specifications are ASHRAE [16], NEBS [12], and ETSI [17]. A summary of these environmental classes is given below in Tables 5, 6, 7, and 8.

Table 5 Summary of ASHRAE 2011 Thermal Guideline Classes [16]

Classes (a)	Equipment Environmental Specifications							
	Product Operations (b)(c)					Product Power Off (c) (d)		
	Dry-Bulb Temperature (˚C) (e) (g)	HumidityRange, non-Condensing (h) (i)	Maximum Dew Point (˚C)	Maximum Elevation (m)	Maximum Rate of Change(˚C/hr) (f)	Dry-Bulb Temperature (˚C)	Relative Humidity (%)	Maximum Dew Point (˚C)
Recommended (applies to all A classes, evaluate ITE metrics in this paper for conditions outside this range)								
A1 to A4	18 to 27	5.5ºC DP to 60% RH and 15ºC DP						
Allowable								
A1	15 to 32	20% to 80% RH	17	3050	5/20	5 to 45	8 to 80	27
A2	10 to 35	20% to 80% RH	21	3050	5/20	5 to 45	8 to 80	27
A3	5 to 40	-12˚C DP & 8% RH to 85% RH	24	3050	5/20	5 to 45	8 to 85	27
A4	5 to 45	-12˚C DP & 8% RH to 90% RH	24	3050	5/20	5 to 45	8 to 90	27
B	5 to 35	8% RH to 80% RH	28	3050	NA	5 to 45	8 to 80	29
C	5 to 40	8% RH to 80% RH	28	3050	NA	5 to 45	8 to 80	29

Table 6 Summary of NEBS equipment aisle[1] air temperature and humidity limits [12].

Conditions	Limits
Temperature • Operating (up to 1829m [6000ft]) • Short-term[2] • Short-term with fan failure	5 to 40°C -5 to 50°C -5 to 40°C
Rates of Temperature Change • Operating	30°C/hour
Relative Humidity • Operating • Short-term[2]	5 to 85% 5 to 93% but not to exceed 0.026 kg water/kg of dry air

Notes:

1. Equipment aisle refers to conditions at a location 1524 mm (60 in) above the floor and 381 mm (15.0 in) in front of the equipment. Equipment test temperatures are defined in

Section 5.1, “Temperature, Humidity, and Altitude Test Methods,” based on equipment configuration (frame-level or shelf-level) and air-inlet location.

2. Short-term refers to a period of not greater than 96 consecutive hours, and a total of greater than 15 days in 1 year. (This refers to a total of 360 hours in any given year, but not greater than 15 occurrences during that 1-year period.)

Table 7 Summary of ETSI Class 3.1 and 3.1e environmental requirements [17]

	Continuous Operation	Class 3.1: ≤10% of Operational Hours	Class 3.1e: ≤1% of Operational Hours
Temperature Ranges	10 to 35°C	5 to 10°C, and 35 to 40°C	-5 to 5°C, and 40 to 45°C
Humidity Ranges	10% to 80%RH[1]	5 to 10%RH[2], and 80 to 85%RH[3]	5 to 10%RH[2], and 85 to 90%RH[3]

[1]With minimum absolute humidity of no less than 1.5 g/m^3 and a maximum absolute humidity of no more than 20g/m^3

[2]With minimum absolute humidity of no less than 1 g/m^3.

[3]With maximum absolute humidity of no more than 25 g/m^3.

Note: maximum rate of temperature change for continuous operation, Class 3.1 and Class 3.1e is 0.5°C/minute averaged over 5 minutes.

Table 8 Summary of maximum altitude ratings and de-rating values.

Specification & Class	Minimum Altitude	Maximum Altitude	De-Rating (Metric Units)	De-Rating (I-P Units)
ASHRAE - Class A1	NA	3050m (10000ft)	1°C/300m above 950m	1.8°F/984ft above 3117ft
ASHRAE – Class A2	NA	3050m (10000ft)	1°C/300m above 950m	1.8°F/984ft above 3117ft
ASHRAE – Class A3	NA	3050m (10000ft)	1°C/175m above 950m	1.8°F/574ft above 3117ft
ASHRAE – Class A4	NA	3050m (10000ft)	1°C/125m above 950m	1.8°F/410ft above 3117ft
NEBS GR-63-CORE	NA	3960m (13000ft)	1°C/305m above 1829m	2.2°F/1000ft above 6000ft
ETSI – Classes 3.1, 3.1e	70kPa (approx. 381m (1250ft) below sea level	106kPa (approx. 3048m (10000ft) above sea level	NA	NA

If one compares the temperature rating survey in Figure 5 to the standards above, it is apparent that many networking products have supplier-defined temperature ratings that exceed the commonly referenced industry standards. A likely explanation for this is the industry standards were written for nominal air supply temperatures and weren’t written to include micro-environments and heat build-up inside enclosed racks. In Figure 17 below the server exhaust temperatures from Figure 13 are superimposed on the bubble chart survey of maximum temperature ratings.

Fig. 17 Server exhaust temperature range super-imposed on benchmarking data.

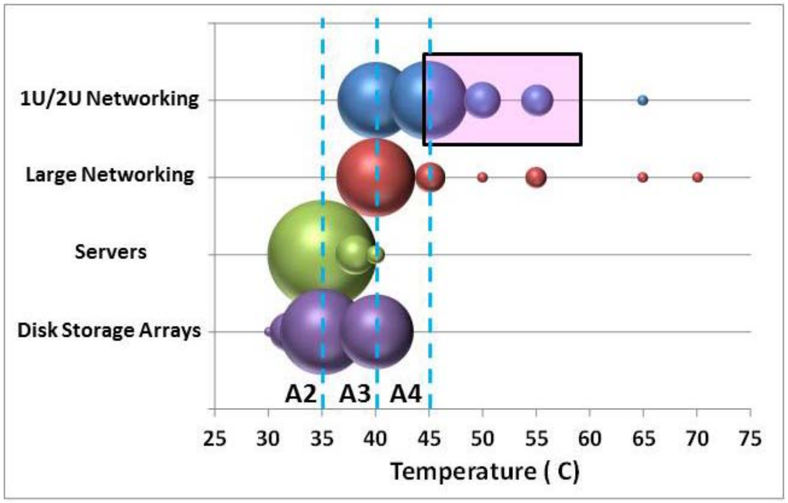

According to our benchmarking survey results, nearly all of the networking equipment has temperature specs that meet ASHRAE Class A3 and NEBS, with some that meet and exceed Class A4 specifications; but ASHRAE also defines the location of the measurement of that temperature to be in the "front of the cabinet or rack" and "50 mm (2 in.) from the front of the equipment". It is implied that the equipment inlet surface is within 2 inches of the defined measurement location which is only the case for networking equipment when it is mounted in the front of the rack. For rear mounted networking equipment that does take its cooling air directly from the data center cold air stream, the Class A3, NEBS, and A4 ratings are probably sufficient even though the equipment inlet surface is much farther removed.

However, if a small form factor (e.g. 1U/2U) switch was installed in a non-ideal, but common, installation such as behind a blanking panel, the equipment would probably need to have a much higher maximum temperature rating such as indicated by the server exhaust temperature range shown above in Figure 17. In the case of a blanking panel or enclosed environment, the temperature rise generated by the switch itself can also add to the inlet air temperature rating needed (see Figure 16). This explains why it is possible, under some circumstances, for a top of rack network switch that is rated at 40 or 45°C air inlet temperature to overheat and possibly lose functionality.

The growing adoption of higher data center temperatures and economization may increase the probability of a situation arising where some of the data center networking equipment exceeds its maximum rated temperature. Exceeding the manufacturer's maximum temperature rating is not allowed. This situation can cause a partial loss of functionality with compromised data integrity or even complete loss of functionality. Non-ideal installations such as blanking panels and side to side cooling architectures may not be sustainable going forward for small form factor networking products. Thus, ASHRAE TC9.9 recommends networking equipment that is intended to be installed in a confined space where heat build-up could occur be designed to a higher maximum temperature rating than equipment designed to take its cooling air directly from the data center cooling air stream such as at the front face of the rack.

The telecom industry is having similar thermal challenges with their networking equipment and they are actively addressing the problem. In the latest revision of GR-63-CORE Issue 4 [12], Telcordia issues a strong recommendation in favor of front to rear cooling architectures:

> *"Air-cooled equipment shall utilize only a rear-aisle exhaust. Air-cooled equipment should utilize a front-aisle inlet. Air-cooled equipment with other than front-aisle to rear-aisle airflow should be fitted with manufacturer provided air baffles/deflectors that effectively re-route the air to provide front-aisle to rear-aisle airflow. Equipment requiring air baffles/deflectors for airflow compliance shall be tested by the manufacturer for compliance to GR-63-CORE with such hardware in place".*

GR-63-CORE does still allow side to side and non-front to rear air flow architectures, however, those pieces of equipment must be designed and tested to 60°C for equipment that takes up the entire rack (frame level) products and 65°C for smaller stand-alone equipment that fits inside the rack (shelf level) products. Coincidentally these maximum temperature ratings are very similar to the maximum temperature ratings indicated in the previous discussion for small form factor networking equipment in a non-ideal installation in a data center. From the information in GR-63-CORE, it is clear the telecom industry is trying to standardize networking equipment on a front to rear air flow and, any equipment that can't be designed or adapted for front to rear air flow will have to be designed and tested to a higher operational temperature.

6 Reliability

Reliability is an important consideration for networking equipment because networking is one of the critical equipment types that determine data center reliability and uptime. The networking space has more models and more different equipment designs, more installation options, and perhaps an even wider spectrum of end use applications than other IT equipment such as servers and storage arrays. In addition, some networking equipment contains unique technology such as lasers. Given the breadth of the networking equipment application space, this paper will focus on general reliability guidelines and trends for data center networking equipment. While a significant amount of discussion below is devoted to reliability, the user is more likely to experience a temporary loss of functionality from over-heating than from a hard reliability type of equipment failure.

Chapter 2 of the ASHRAE Thermal Guidelines for Data Processing Environments book [15] provided general guidelines for the reliability of data center IT equipment, including networking, as a function of temperature (see Figure 18 below). The thermal guidelines book data only went as high as 45°C so an extrapolated estimate (dashed line) has been added to project IT equipment reliability to a higher inlet temperature range that may be encountered by some networking equipment. We believe the extrapolation using a gradual increase in failure rate is reasonable based on published papers [18, 19, 20] and anecdotal data from member companies.

Figure 18 Failure rate of IT equipment as a function of temperature [15].

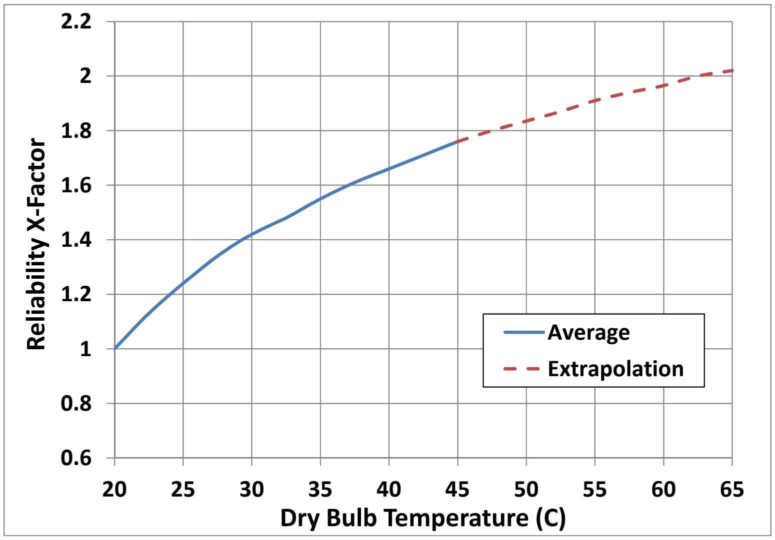

The y-axis is a dimensionless reliability X-factor that represents the ratio of failure rate at a given dry bulb temperature to the failure rate at 20°C. This curve is only valid for air inlet temperatures at or below the maximum rating of a particular piece of equipment. The curve above is not intended to imply operation above the manufacturer's specified maximum air inlet temperature. Per the benchmarking data in Figure 5, commercially available networking products are available with inlet air temperature ratings as high as 70°C.

A quick inspection of the plot shows the relative failure rate increases by about 1.8X from 20 to 45°C. This is significantly less than the 2X per 10°C traditional reliability prediction models (e.g. MIL-217, Telcordia) would suggest. The data in Figure 18 was derived from a large population consisting of different types of IT equipment which was deployed in real-world applications over an extended period of time.

Using the failure rate data from Fig. 18 one can compare and contrast the failure rate impact of ideal versus non-ideal networking equipment installations. In an ideal equipment design and installation, the equipment takes its inlet cooling air directly from the cold aisle supply air provided by the data center. However, in a non-ideal design and installation such as a top of rack switch installed behind a blanking panel, the equipment inlet air temperature is significantly hotter because a micro-environment is formed that traps heat. Consider a data center with a 25°C cold aisle air temperature. In this scenario, the ideal equipment designs and installations would have a 25°C air inlet temperature. A top of rack switch behind a blanking panel would see the temperature of the server exhaust plus 5 or 10°C of temperature rise from the heat exhaust of the switch itself. For a cold aisle temperature of 25°C, Fig. 13 shows a server exhaust temperature in the range of 50°C. With a 5°C temperature rise from the switch, this becomes an air inlet temperature for the switch of approximately 55°C. The failure rate comparison is shown below in Fig. 19.

Fig. 19 X-factor failure rate comparison of servers and networking equipment.

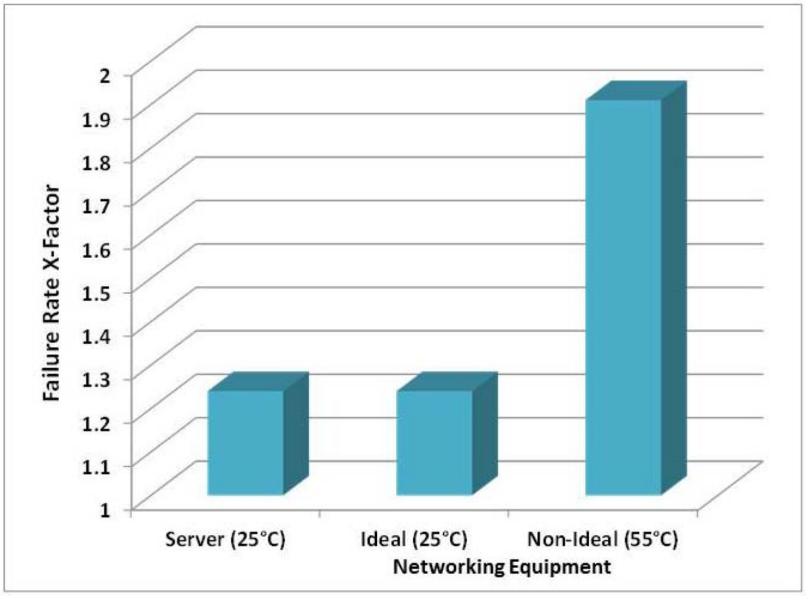

The y-axis of the plot is IT equipment failure rate relative to a 20°C air inlet temperature. The ideal and non-ideal designs are networking equipment and the non-ideal design x-factor is a top of rack switch behind a blanking panel.

For servers and ideal networking equipment designs and installations the X-factor is 1.24 (relative to equipment operated at 20°C) whereas for the non-ideal networking equipment it is 1.91 (relative to equipment operated at 20°C). Thus, the additional 30°C of temperature rise from the micro-environment will cause the top of rack network switch to have a 54% higher failure rate. A medium or large data center that uses the topology shown in Fig. 4 could have quite a few of these top of rack networking switches. And, if these top of rack switches were installed behind blanking panels, i.e. in a non-ideal manner, a large and important segment of the critical IT equipment in the data center would be subject to a 54% failure rate penalty. This would seem to be out of step with good data center reliability and availability design practices. Also, a pre-heated micro-environment can also cause the equipment to overheat exceeding its maximum temperature rating and losing functionality and possibly compromising data integrity. The 55°C air inlet temperature in our scenario exceeds the maximum temperature rating of more than 90% of commercially available top of rack network switches.

There is an additional concern with micro-environments – if the networking equipment is being operated in a temperature range well above what it was designed for, the user could encounter wear-out failure mechanisms that cause the equipment to fail well in advance of its intended design lifespan. These wear-out failure mechanisms include premature loss of electrolyte in the electrolytic capacitors commonly used in power supplies, and silicon wear-out mechanisms such as electro-migration (EM) and time dependent dielectric breakdown (TDDB). Operating networking equipment, or any IT equipment for that matter, in a temperature range above its design point is risky and should be avoided.

One potential solution to the higher failure rate caused by a micro-environment is to mitigate it by designing the networking equipment to operate in a higher temperature environment. The equipment could be designed and built with components that have higher temperature ratings. And, the equipment could be designed with larger air movers that would move a larger volume of air through the equipment and remove more heat reducing the temperature rise in the components. However, both of these approaches have drawbacks. It is difficult and expensive to find components temperature ratings that will support a 55°C or higher system inlet temperature. While larger air movers are a potential solution they would consume significantly more power. The power consumption of these devices goes up as the cube of the RPM. From a reliability and cost perspective, the best solution is simply to provide cooling air from the cold aisle face of the rack directly to the networking equipment as shown earlier in the section on ideal networking equipment designs and installations. This approach is consistent with recommendations in Telcordia GR-63-CORE issue 4 [12].

The reliability and failure rate of networking equipment is influenced by other factors besides temperature including humidity, dust, and gaseous pollution. These factors are particularly important for economized data centers, data centers located in polluted regions of the world, and data centers installed in specialized environments. The impact of these factors on hardware, recommended limits and recommended best practices are detailed in the 2011 Gaseous and Particulate Contamination Guidelines for Data Centers white paper [21] and the ASHRAE Particulate and Gaseous Contamination in Datacom Environments book [22].

7 Practical Installation Considerations

As was previously stated, an ideal installation creates an acceptable environment for the switch without disturbing the integrity of the cooling strategy for the rest of the rack. If a rear mount switch doesn’t have a high enough maximum air inlet temperature to operate behind a blanking panel, it will need to be fitted with a duct kit to allow it to take inlet air from the cold aisle at the front face of the rack. However, most switches don’t come with an accessory duct kit. The end user often has to investigate options on his own.

One of the easiest decisions to make is where to mount a top of rack switch - either in front or in rear of the rack. If the equipment is mounted in front so the ports face forward, there is likely a need for a pass through to the rear of the rack for cabling. Commercially available products exist that allow the cables to pass through some type of grommet to block airflow. If the grommet is of the brush type, it is advisable to have more than one row of brushes. Single row grommets typically do not do a very good job of blocking airflow.

For a currently available or legacy switch mounted in the rear of a rack, there are a number of decisions to be made, and there really is no such thing as a “best practice”. Sometimes trial and error yields the best results. The most straight forward strategy is to install solid blanks in front of the switch and force it to consume rear rack airflow for cooling. This is air that has been preheated by other IT equipment and will be very close to IT exhaust temperature. Regardless of the airflow direction, the inlet temperature to the switch could be quite high. It is advisable to measure rear rack temperatures before installing a switch to make sure temperature rating of the networking equipment is compatible with the air temperature in the rack. If measurements aren’t possible, some IT vendors provide delta-T estimates in their on-line power calculators. Most commercially available switches are rated for a maximum inlet air temperature of between 40°C and 50°C. Also, the presence of a perforated rear door can have an effect. Without the rear door, momentum often carries server exhaust away, out of the rack. A rear (R) inlet switch may have access to cooler hot aisle air than when door is present especially if the front of the rack is blanked. A rear (R) cooling air inlet switch can be a better choice in this situation than a side or front inlet.

Even though blanking in front of the switch might be considered a best practice, in *some* instances, leaving a gap will lower the switch temperature without significantly affecting the other IT equipment. This is more common on racks without front and rear doors because the pressure differential between the front and the rear of the rack is not as great.

Elevation of the switch is definitely a variable to consider. Most switches are mounted in the top of the rack. The top of the rack tends to be warmer than the middle or bottom of the rack, especially with raised floor cooling.

If the switch is side-to-side cooled, you may need to consider blocking lateral air flow between adjacent racks. When using open frame style racks with no cabinet, side-to-side switches in adjacent racks can set up a situation where the hot exhaust from one switch feeds directly in to the cooling inlet of the switch in the adjacent row. This situation should be addressed by installing an air flow baffle between the two racks to block rack to rack air flow.

Rather than relying on rear rack temperature air to cool the switch, some prefer to use ductwork to gain access to cool air from the front of the rack. Unless the switch has been over-designed to accommodate external pressure, the introduction of ductwork adds some uncertainty. The ductwork becomes an extra pressure drop the switch has to work against. The resulting reduction in airflow, if that happens, may raise internal temperatures and lower the "inlet" temperature to which the switch issues a thermal warning. Testing of a front-to-rear and side-to-side switch produced flow reduction values of 10% for the straight duct with linear flow and a full 1/3 reduction for the right angle turning flow pattern associated with the side-to-side switch. This reduction in air flow can also be exacerbated by racks with passive containment like a chimney, rear door coil, or any other tightly ducted passive containment system. Besides the resistance of pulling air through the ductwork, the switch may also be pushing against increased resistance of the contained space. The resistance of the contained space is also variable in nature and related to the amount of air moving through the rack, so it is hard to give definitive advice in this instance. If a switch with a high enough maximum air inlet temperature isn't available that can work with the air flow resistance, one might have to consider an active cooling solution (e.g. fan assisted) rather than a passive solution such as duct work alone.

There are also commercially available active ducts such as the one shown below in Fig. 20 to help draw air from the front of the rack and route it to the switch inlet. These don't always make airtight connection with the switch, and the inlet temperature for the switch is a mix of front and rear rack temperature. These active products do provide a solution if the switches aren't strong enough to pull air through a passive duct. When considering an active solution, make sure the duct fans aren't undersized for the switch flow rate, and make sure the air is moving in the proper direction. There are some simple fan panels on the market that you can place in front of the switch, but some, by default, move the air in the wrong direction (toward the front of the rack).

Fig. 20 Example of commercially available fan assisted duct work.

Ducting Inside the Rack

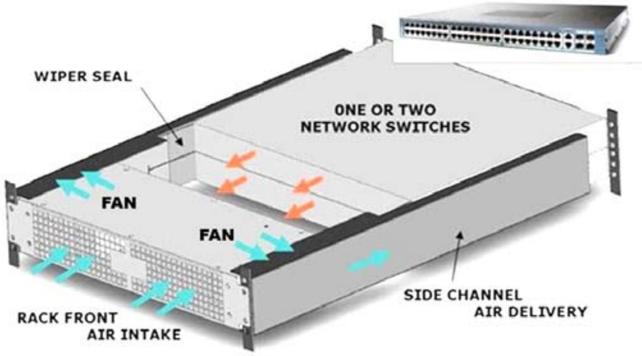

Grill on Front of Rack

8 ASHRAE TC9.9 Recommendations

Since the first ASHRAE recommended data center environmental conditions were published in 2005, there have been significant changes in data center operating conditions. The latest ASHRAE thermal guidelines [15] have an A4 range that allows a 45°C maximum inlet temperature and new data center cooling technologies such as economization are rapidly gaining acceptance. Changes to air flow and rack level design, such as aisle-level containment ("hot aisle, cold aisle"), are becoming commonplace. The designs of most types of IT equipment have undergone changes to reduce energy consumption. One of the most common changes is reducing the speed of the cooling fans – this has the effect of raising the IT equipment exhaust temperature. Future designs and installations of networking equipment will need to consider all of these factors and successfully integrate with other types of IT equipment into a rack level solution with little or no re-engineering on the part of the end user.

The data center and IT equipment landscape of the future will require a new set of thermal design recommendations for networking equipment:

- **Use equipment with front of rack to rear air flow** - Adopt front of rack to rear of rack air flow wherever possible for new equipment designs and installations to give the IT equipment direct access to the data center cold air supply.
- **Draw cooling air flow from the front face of the rack** - instead of taking it from inside the rack where mixing of hot and cold air streams can occur. In the case of special rack designs, such as designs that take cooling air from the sub-floor, the same principles apply although the implementation may be slightly different – the networking equipment needs to have direct access to the data center cold air supply.
- **Use duct work to draw cooling air from front face of rack**. For new designs of small top of rack switches, ASHRAE TC9.9 recommends IT equipment manufacturers provide duct work ("snorkel kits") that extends the depth of the chassis to the front of the rack so the equipment can draw its cooling air from the front panel of the rack if the switch is

mounted to the rear of the rack, as is common. The networking equipment needs to be designed and qualified to function with the duct work over its entire specified temperature range.

- **Non-front to rear air flow cooling designs are discouraged.** If networking equipment must be designed with non-front to rear air flow, it needs to be designed to a higher air inlet temperature so it can handle all of the common types of data center installations without over-heating. It should be noted that, with careful design and installation practices, side to side air flow equipment can be made to work.
- **Design front to rear cooled networking equipment to a minimum of ASHRAE Class A3 (40°C), preferably Class A4 (45°C).**
- **Use duct work to convert side to side switches to front to rear cooling.** For side to side air flow designs of small top of rack switches, users should consider installing duct work ("snorkel kits") to convert existing equipment to front of rack to rear of rack air flow so the cooling air is taken from the front face of the rack. The fans in the switch should have enough power to overcome added air flow resistance of the duct work and the addition of the duct work shouldn't significantly raise key component temperatures inside the equipment.
- **Exhaust hot air at the rear face of the rack –** this practice will prevent mixing of the exhaust air stream with the cold air stream. In the case of racks with non-front to rear air flow, the implementation may be slightly different.
- **Isolate hot and cold air streams from one another inside the rack –** prevent mixing of hot and cold air streams. Ensure there are no bypass paths around the networking equipment or through the rack to allow cold air to mix with hot air.
- **Design simple, preferably linear, cooling air flow paths** - design of cooling air flow paths inside the equipment and inside racks needs to be kept as simple as possible, preferably linear with no bends. Where linear front to rear air flow is not possible, the equipment design needs to account for the added air flow impedance of any bends.
- **Provide detailed thermal and air flow installation instructions that comply with ASHRAE TC9.9 best practices –** this will ensure the cooling air to the equipment inlet stays within specification ratings at all times. The usage instructions should specify which types of installations the equipment is designed to support. IT equipment users and data center operators must carefully follow manufacturer's instructions and ensure the manner in which they have installed the equipment provides a steady stream of cooling air that is within the equipment manufacturer's specification limits at all times.
- **Provide enough cabinet width for adequate cooling air flow, especially for side to side cooled equipment** - Where side to side cooled equipment must be used, ensure there is adequate width inside the cabinet to accommodate the cooling air flow needed with all of the cabling in place. Networking equipment is cabling intensive and large amounts of cabling can impede air flow. Additional cabinet width may also be needed to accommodate the duct work needed to convert side to side air flow equipment to front to rear.
- **Deliver cooling air uniformly along the width/height of the equipment to ensure uniform cooling.** Some equipment designs require a uniform supply of cooling air along the length of the air inlet while others may be designed to work better with a non-uniform flow of cooling air. One should understand the thermal design of the equipment and engineer the cooling air distribution to the air inlet in a way that works with the

equipment design and provides the correct amount of cooling internal to the IT equipment.

- **Use partitions to prevent hot exhaust from feeding the cold air inlet of another piece of equipment**. Install partitions between adjacent open 2 and 4 post style racks that contain side to side cooled rack mounted equipment to prevent one piece of equipment from exhausting hot air into the cold air inlet of another.
- **Make hot-aisle touch points out of non-metallic materials** - consider making touch points such as power supply handles from non-metallic materials such as plastic and rubber (instead of metal) that have higher touch point temperature limits and will be more comfortable for the user to hold.
- **Maintain large enough cable radii so the cabling doesn't impede air flow.**
- **Seal cable feed-throughs to improve data center cooling efficiency**. Seal feed-throughs where networking equipment cables pass through an air flow bulk head such as a raised floor. Cold air leakage through cable ports can reduce data center cooling efficiency and increase operational costs. For example, brushes can be used to seal gaps and still allow for the pass through of cables. Patch panels, with connectors on both sides of the panel, can be as a means of passing cables through an air-tight bulkhead.
- **Comply with recommendations in 2011 Gaseous and Particulate Contamination Guides for Data Centers**. The air stream in the data center environment should be free of gaseous pollutants and dust to the levels recommended in the 2011 Gaseous and Particulate Contamination Guidelines for Data Centers white paper [21]. If the networking equipment has dust filters, these should be inspected and cleaned regularly at the interval recommended by the manufacturer. The networking equipment usage environment should be free of lint as lint can clog filters and air inlets and can compromise the cooling ability of the equipment.

9 Summary

The topology of networking equipment in a typical data center includes core, distribution and edge equipment. This equipment has a wide variety of chassis sizes and cooling air flow designs as well as different rack installation configurations. Over the last several years there have been significant changes to data center operating environmental conditions that impact networking equipment design and installation. Some of these changes include: a) use of higher data center operating temperatures to save energy on cooling costs, b) growing adoption of economization, c) wider allowable operating temperature ranges, and d) energy saving reduction of fan speeds internal to the IT equipment that are causing IT equipment exhaust temperatures to rise. Given the backdrop of these changes, networking equipment design and installation practices should be revised to avoid potential problems with equipment over-heating. Network equipment that exceeds its maximum temperature rating may experience a temporary loss of functionality and it is possible this type of event could compromise data integrity.

A networking equipment design that conforms to the ASHRAE TC9.9 recommendations would be front to rear cooled, with cooling air taken from the front face of the rack, and hot air exhausted at the rear face of the rack. Networking equipment that may need to be installed in an enclosed space where heat build-up could occur should consider designing for higher air inlet temperatures the equipment may encounter. The common practice of locating 40 or 45°C rated equipment behind a blanking panel is risky and is probably not sustainable.

If new networking products and installations follow the ASHRAE TC9.9 design and installation recommendations, concerns of over-heating will be largely removed, equipment failure rates will be reduced, and overall data center uptime should improve. The ASHRAE TC9.9 recommendations will also simplify the job of the end user by making it easier to integrate the networking equipment into a rack solution with other types of IT equipment.

10 References

[1] Best Practices for Datacom Facility Energy Efficiency, Second Edition, ASHRAE, 2009, 230 pp., ISBN/ISSN: 978-1-933742-27-4.

[2] Liquid Cooling Guidelines for Datacom Equipment Centers, ASHRAE, 2006, 100 pp., ISBN/ISSN: ISBN-10: 1-933742-05.

[3] Design Considerations for Datacom Equipment Centers, Second Edition, ASHRAE, 2009, 208 pp., ISBN/ISSN: 1-931862-94-X.

[4] Datacom Equipment Power Trends and Cooling Applications, 2nd Edition, ASHRAE, 2012, 144 pp., ISBN/ISSN: 978-1-936504-28-2.

[5] "Networking Vendors Can't Take the Heat", Data Center Dynamics, November 14, 2012; http://www.datacenterdynamics.com/focus/archive/2012/11/networking-vendors-cant-take-heat-facebook-says

[6] ISO/IEC 7498-1, Information technology – Open Systems Interconnection – Basic Reference Model: The Basic Model.

[7] The Essential Guide to Telecommunications, 5th edition, Annabel Z. Dodd, Prentice-Hall, 2012, 483 pp.

[8] Telecommunications Essentials, 2nd edition, Lillian Golleniewski, Addison-Wesley, 2009.

[9] Networking Bible, Barry Sosinsky, Wiley Publishing, 2009, 890 pp.

[10] Networking, 2nd edition, Jeffrey S. Beasley, Pearson Education Inc., 2009, 695 pp.

[11] Electronic Industries Alliance EIA-310-D, Cabinets, Racks, Panels, and Associated Equipment.

[12] NEBS Requirements: Physical Protection, GR-63-CORE Issue 4, April 2012, Telcordia Technologies, Inc.

[13] 2008 ASHRAE Environmental Guidelines for Datacom Equipment, ASHRAE white paper.

[14] Telcordia GR-3028-CORE: Thermal Management in Telecommunications. Central Offices, issue 1, December 2001.

[15] Datacom Equipment Power Trends and Cooling Applications, 2nd edition, ASHRAE, 2012, 144 pp.

[16] Thermal Guidelines for Data Processing Environments, 3rd edition, ASHRAE, 2012, 150 pp.

[17] European Telecommunication Standard - Equipment Engineering (EE); Environmental conditions and environmental tests for telecommunications equipment, Part 1-3: Classification of environmental conditions, stationary use at weather protected locations, ETS 300 019-1-3.

[18] N. El-Sayed, I. Stefanovici, G. Amvrosiadis, A.A. Hwang, and B. Schroeder, “Temperature Management in Data Centers: Why Some (Might) Like It Hot”, proceedings of Thirteenth InterSociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, ITHERM 2012, San Diego CA, USA.
[19] B. Schroeder, E. Pinheiro, W-D. Weber, “DRAM Errors in the Wild: A Large Scale Field Study”, SIGMETRICS/Performance ’09, June 15-19, 2009 Seattle WA, USA.
[20] E. Pinheiro, W-D. Weber, L.A. Barroso, “Failure Trends in a Large Disk Drive Population”, 5th USENIX Conference on File and Storage Technologies, (FAST’07), February 2007.
[21] 2011 Gaseous and Particulate Contamination Guidelines for Data Centers, white paper published by the ASHRAE Technical Committee 9.9 Mission Critical Facilities, Technology Spaces and Electronic Equipment.
[22] Particulate and Gaseous Contamination in Datacom Environments, ASHRAE, 2009, 94 pp.
[23] Sanjay K. Roy, “An Equation for Estimating the Maximum Allowable Surface Temperatures of Electronic Equipment”, 27th IEEE SEMI-THERM Symposium, 2011, p. 54.

APPENDIX A – Definition of Acronyms and Key Terms

Acronym/Term	Definition
Air-side economization	The use of outside air for data center cooling instead of using a chiller or air conditioning. Air-side economization is typically used only when outside air conditions are within a defined acceptable range.
ANSI	American National Standards Institute
ASIC	Application Specific Integrated Circuit
ASTM	American Society for Testing and Materials
ATIS	Alliance for Telecommunications Industry Solutions
Blanking panel	A solid panel typically installed on the front face of the rack to prevent air flow from leaking through the rack
Bridge	A device that connects multiple network segments along the data link layer.
Cold aisle	In a data center where hot and cold air streams are contained inside the rows between racks, the cold aisle is the aisle the contains the cold air stream.
Core Network	The core network function is a gateway through which all data entering and exiting the data center must pass. The core network equipment is connected either directly or indirectly to every device in the data center. The core switch is also connected to a service provider (SP), which is the “pipe” through which all data into and out from the data center passes and connects to the internet.
CPE	Customer premise equipment
Data center	A building or portion of a building whose primary function is to house a computer room and its support areas; data centers

	typically contain high-end servers and storage products with mission-critical functions. A building or dedicated space(s) within a larger facility whose primary function is to house a computer room(s) and its support areas; data centers typically contain high end servers & storage components and related communications & connectivity hardware providing mission critical functions
Dew Point	The dew point is the temperature below which the water vapor in a volume of humid air, at a given constant barometric pressure, will condense into liquid water at the same rate at which it evaporates. Condensed water is called dew when it forms on a solid surface.
Distribution Network	Acts as an intermediate level between the core and edge levels of the network and it offloads some of the work the core network equipment needs to do. The distribution level is useful for passing data between machines inside the data center, or aggregating ports to reduce the number of physical ports required at the core.
Dry bulb	The temperature of air indicated by a thermometer
DSL	Digital Subscriber Line
ECMA	ECMA International – European association for standardizing information and communication systems. The organization was formerly known as “European Computer Manufacturers Association” and the trademark “ECMA” was kept for historical reasons.
Edge Network	Consists of switches that connect directly to devices that are generating or consuming the data, and then passing the data up to the distribution or core.
EM	Electro-migration, a wear-out mechanism of the metal interconnects in silicon integrated circuits
Equipment (IT)	Refers to, but not limited to, servers, storage products, workstations, personal computers, and transportable computers; may also be referred to as electronic equipment or IT equipment.
ETSI	European Telecommunications Standards Institute
Form factor	The height of the equipment chassis measured in units of “U” where 1U = 1.75 inches
Fresh air cooling	The use of outside air for data center cooling. Another term for this is air-side economization
Gateway	A device that is placed at a network node and interfaces with another network that uses different protocols.
Hot aisle	In a data center where the hot and cold air streams are contained between rows of racks, this is the aisle that contains the hot exhaust air from the IT equipment
Hub	A device that connects multiple Ethernet segments, making

	them act as a single segment.
HVAC	Heating, Ventilating and Air Conditioning
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization
IT	Information Technology
LAN	Local Area Network
NEBS	Network Equipment Building Standard
Networking rack	A standard 19 inch rack in an enclosed cabinet which is approximately 32 inches wide
OSHA	Occupational Safety and Health Administration (US)
OSI	Acronym for the ISO/IEC 7498-1 standard Open Systems Interconnection (OSI) 7-layer model for computer network protocols
PFC	Acronym for Priority Flow Control (PFC) rules. PFC rules prioritize which data is sent first and which data can wait.
PHY	PHY is an abbreviation for the physical layer of the OSI model
Rack	Frame for housing electronic equipment
Relative humidity	a) ratio of the partial pressure or density of water vapor to the saturation pressure or density, respectively, at the same dry-bulb temperature and barometric pressure of the ambient air, b) ratio of the mole fraction of water vapor to the mole fraction of water vapor saturated at the same temperature and barometric pressure; at 100% relative humidity, the dry-bulb, wet-bulb, and dew-point temperatures are equal.
Repeater	A device that connects multiple Ethernet segments, making them act as a single segment.
Router	A specialized network device that determines the next network point to which it can forward a data packet towards the final destination of the packet
RPM	Revolutions Per Minute
Server rack	A standard EIA 19 inch rack in an enclosed cabinet which is approximately 24 inches wide
SOHO	Small Office, Home Office
SP	Service Provider
Switch	A device that allocates traffic from one network segment to certain lines which connect the segment to another network segment. Unlike a hub, a switch splits the network traffic and sends it to different destinations rather than to all systems on the network.
TDDB	Time Dependent Dielectric Breakdown - a gate oxide wear-out failure mechanism for silicon integrated circuits
Telecom	Abbreviated term for telecommunications
TOR	Top Of Rack

URL	Acronym for Uniform Resource Locator, abbreviated URL, also known as web address, which is a specific character string that constitutes a reference to a resource
Water-side economization	The use of outdoor air to accomplish data center cooling where the means of heat transfer is water instead of air.
X-factor	A dimensionless metric that measures the relative hardware failure rate at a given constant equipment inlet dry-bulb temperature when compared to a baseline of the average hardware failure rate at a constant equipment inlet dry-bulb temperature of 20°C (68°F).

APPENDIX B – Acoustics

Overall acoustic levels in networking equipment have risen steadily over time. Compounding rising acoustic levels are energy saving trends such as increasing data center operating temperatures and the growing adoption of air and water-side economization. Higher operating temperatures, both at the data center level and at the IT equipment air inlet, require the movement of more air to remove the same amount of heat. Empirically, it has been determined that fan sound power increases with fan speed according to the following relationship:

$$Lw_2 - Lw_1 = 50 \log_{10}(n_2/n_1)$$

where Lw = Sound power level (dB), and n = fan speed (RPM). This relationship is shown graphically in Figure 21.

Fig. 21 Example of fan sound power as a function of fan speed.

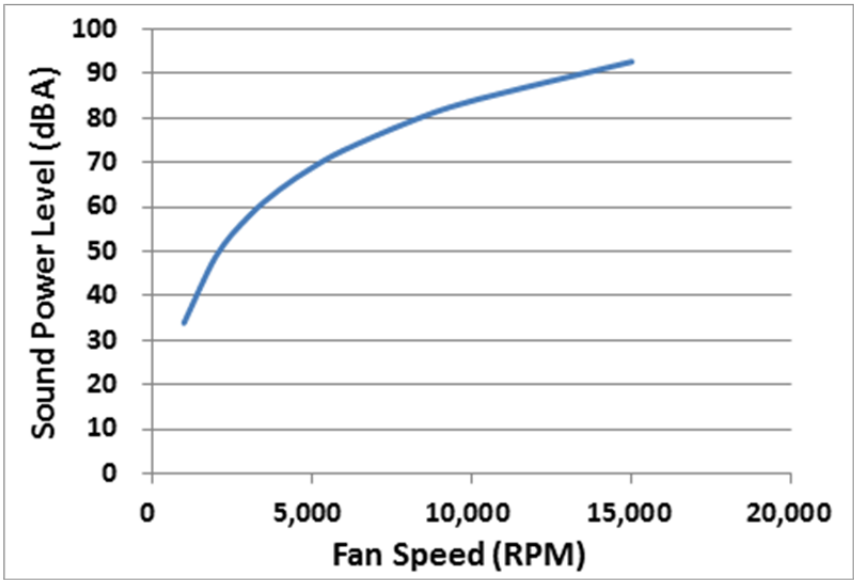

As the networking equipment inlet air temperature increases, the fan speed will increase causing a logarithmic rise in the sound power level. The exact relationship between inlet air temperature and fan speed is unique to the design of each piece of equipment - some equipment has sophisticated polynomial fan control while simpler equipment often has a fixed fan speed or rudimentary speed control. An example of the expected increase in A-weighted sound power levels as a function of temperature is shown below in Fig. 22. For equipment with an 85dB(A) sound power at 25°C, this becomes 89.7dB(A) at 30°C. 87dB(A) and 90dB(A) are common regulatory noise thresholds in Europe and the United States, respectively. Thus, in this example, an acoustically compliant data center could become non-compliant with the additional sound power from a 5°C increase in inlet air temperature.

Fig. 22 Example of the expected increase in A-weighted sound power level with temperature.

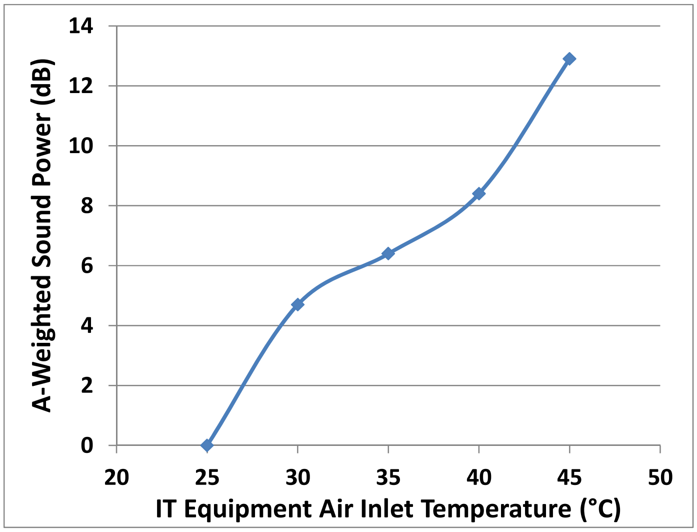

Data for the plot above was taken from Table 2.5 of ASHRAE Thermal Guidelines for Data Processing Environments – 3rd Edition [15].

The most significant acoustic concerns in a data center arise from large end of row and core networking equipment that has a high power density and consumes most or all of a rack. The contribution of small top of rack network switches to overall data center sound power levels is almost negligible.

The design and testing of networking equipment should comply with the testing methods and acoustic criteria listed in Tables 9 and 10. When designing and provisioning a data center, local, state and national acoustic regulations should be carefully consulted – a discussion of these regulations is beyond the scope of this document. In the event prevailing acoustic regulations are not met, the data center operator may need to consider abatement methods such as installing sound absorbing materials on the floor, walls, and ceiling. The interaction between the data center room and the IT equipment can make meeting OSHA and other room-level regulatory requirements fairly complex.

For more detailed information on IT equipment acoustic levels, please consult Section 2.4.2 of the ASHRAE Thermal Guidelines for Data Processing Environments [15].

Table 9 Relevant industry standards for acoustics and acoustic measurements.

Industry Standard Number	Title
ISO 7779; 1999 (ANSI S12.10-2002)	Measurement of Airborne Noise Emitted by Information Technology and Telecommunications Equipment
International Standard ISO 9296:1988 (E)	Acoustics -- Declared noise emission values of computer and business equipment
Statskontoret Technical Standard 26:6	Acoustical Noise Emission of Information Technology Equipment
Standard ECMA-74	Measurement of Airborne Noise Emitted by Information Technology and Telecommunications Equipment. 8th edition (December 2003)
ATIS-0600005.2006 (R2011)	Acoustic Measurement
ETS-300 753	Equipment Engineering (EE); Acoustic noise emitted by telecommunications equipment
NEBS/Telcordia GR-63-CORE	NEBS Requirements: Physical Protection
ISO 3744	Acoustics -- Determination of sound power levels and sound energy levels of noise sources using sound pressure -- Engineering methods for an essentially free field over a reflecting plane
ISO 3745:2012	Acoustics -- Determination of sound power levels and sound energy levels of noise sources using sound pressure -- Precision methods for anechoic rooms and hemi-anechoic rooms

Table 10 Summary of some key room-level acoustic criteria for IT equipment.

Industry Standard Number/Agency	Key Acoustic Criteria
OSHA 29 CFR 1910	"Occupational Safety and Health Standards" - OSHA time limits are 92dbA sound pressure for 8 hours. For anything above 85dbA sound pressure, the employer must provide ear protection. Note: interaction between the room and the IT equipment can make meeting the OSHA requirements complex.
NEBS/Telcordia GR-63-CORE	GR-63-CORE changed to sound power instead of sound pressure. For an attended central office room, the declared sound power limit is 7.8 bels, and 8.3 bels for an unattended space. However, for issue 4 which takes effect in 2013, the requirement for both attended and unattended is 7.8 bels and there is no separate criteria for unattended spaces.
ATIS-0600005.2006 (R2011)	Maximum acoustic noise emission limits for telecommunications equipment to be installed in temperature-controlled environments are: Class 1 Telecommunications equipment room (unattended) – sound power limit of 75 dbA at 27°C Class 2 Telecommunications equipment room (attended) – sound power limit of 75 dbA at 27°C Power room – limit of 87 dbA sound power at 27°C
ETS-300 753	A maximum of 7.5 bels for telecom equipment rooms

Note: to convert dbA values to bels, divide the dbA value by 10.

APPENDIX C – Touch Temperature

In light of data center trends such as increasing power density, higher data center ambient temperatures, and rising exhaust temperatures, the touch temperature of metal, glass and plastic parts is an important consideration for equipment design, rack provisioning and data center design. Two of the widely used industry standards on touch temperatures are summarized below in Tables 11 and 12.

Table 11 Touch temperature limits from Telcordia GR-63-CORE.

Materials	Permitted Temperature (°C) as a Function of Exposure Times	
	Unintentional Contact or Parts Held for Short Periods in Normal Use[1]	Prolonged Use[2]
Metals[3]	55°C	48°C
Non-metals	70°C	48°C

[1]*Short periods are defined as ≤10 seconds.*
[2]*Prolonged use is defined as >10 seconds up to 10 minutes.*

[3] *Metals may be coated, uncoated, plated, and/or have a conversion coating.*

Table 12 Touch temperature limits from IEC60950-1

Object Being Touched	Touch Duration	Metal	Glass and Similar	Plastic and Rubber
Handles	Briefly	60°C	70°C	85°C
Handles	Longer Duration	55°C	65°C	75°C
External surfaces and parts inside that may be touched	NA	70°C	80°C	95°C

Note: The exact duration of the terms “briefly” and “long duration” are not defined in IEC60950-1.

In addition to GR-63-CORE and IEC60950-1, reference [23] compares a wide range of touch temperature standards and offers a predictive method for determining touch temperatures and choosing appropriate materials for electronic products.

Server and storage equipment exhaust temperatures could easily create metal surfaces, such as metal power supply handles, in the range of 55-60°C which would exceed the GR-63-CORE touch criteria for both short period and prolonged use. The IEC60950-1 longer duration criteria would also be exceeded for metal handles. For the design and provisioning of networking equipment inside an enclosed rack space or adjacent to a hot aisle, one may want to consider touch points made of non-metallic materials (e.g. plastic and rubber) which have higher touch temperature limits.

Another consideration is the safety of personnel inside environments with high air temperatures, such as a data center hot aisle. A good discussion of applicable recommendations and regulations can be found in Appendix E of the ASHRAE Thermal Guidelines for Data Processing Environments, 2012, 3rd edition [15].