Ed Eacueo
Data Center Manager

Rack Hygiene

Executive Summary

This paper describes the concept of Rack Hygiene, which positions the rack as an airflow management device, not an inert commodity. It offers solutions at the rack level, by defining a benchmarking methodology for quantifying airflow leakage as a means of setting performance goals for racks.

Today, virtualization is being adopted at an increasing rate. A key driver for the deployment of this technology is the reduction of operating costs associated with the consolidation of server, storage and network devices. By-products of this new virtualized environment include a net reduction in IT equipment and associated space. However, the resultant power and cooling loads are condensed into a smaller footprint and have a dynamic association with the IT processing load.

While a legacy cooling architecture can be adapted to the new environment, it is most often inefficient and/or ineffective due to the following inherent design flaws:

- Hot and cold air mixing
- Misalignment of cooling units and IT racks
- Localized inability to reject heat due to unbalanced heat loads on cooling units
- Excessive distance between cooling units and heat loads
- Air distribution is compromised by excessive cable loading
- Inability of legacy cooling infrastructure to react to dynamic heat loads
- Lack of airflow management to accommodate side to side heat rejection devices
- Oversupply of cold air

Containment strategies are gaining acceptance with IT and facility managers who want to optimize their existing room layouts. As containment of hot and cold air streams increases in popularity, it becomes necessary to seal all air gaps in every critical component in the airflow stream to gain maximum effectiveness. This method is deployed today in well-managed raised-floor data centers, in which sealing all potential air leakage gaps is critical to maintaining uniform, sub-floor, static pressure and airflow distribution.

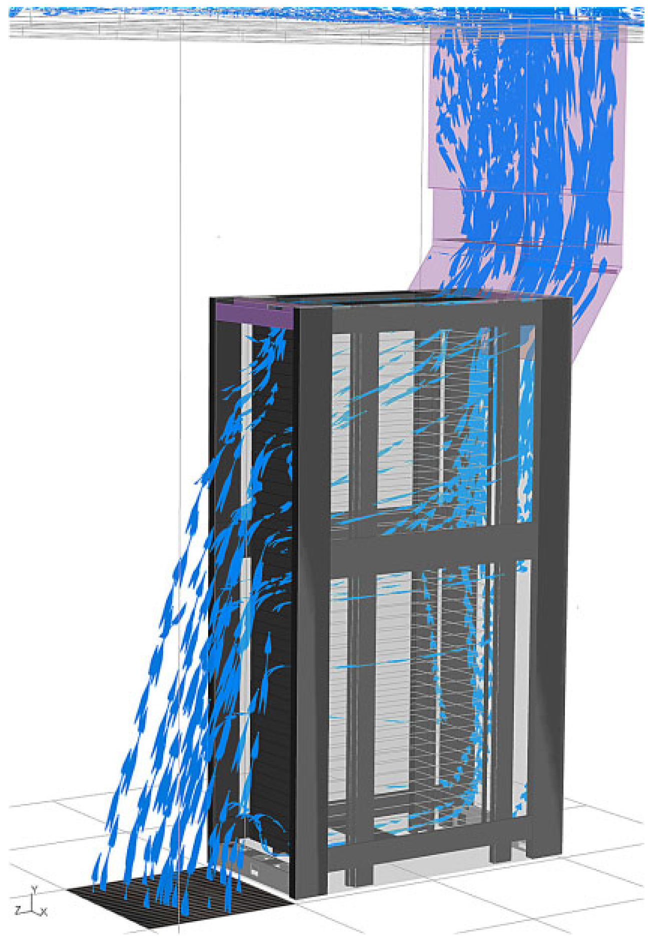

Figure 1. CFD Model of rack using cold air stream to indicate typical air gaps

The “rack” is the forgotten, yet critical, component in the airflow stream. In the rack, air leakage is rampant (Figure 1). It causes recirculation and bypass airflow inefficiencies. Though the EIA-310 industry standard exists for mounting IT equipment in racks, there are no standards for managing hot and cold airflow streams within the rack environment.

Powering Business Worldwide

Rack Evolution

Over the last decade, racks have evolved into a "front-to-back" dominated airflow environment. Glass doors have been replaced with perforated doors which have advanced from 45% open area to today's 65% (and more) open area for maximum airflow. In many instances, front doors are being eliminated altogether.

Roof-mounted fan trays have been replaced with rear-door fans. Cooling coils have been integrated on rear doors to promote "front-to-back" airflow. Heat containment duct systems evolved as a method for rejecting hot air from the rack. Meanwhile, the quantity of power and size of network cables has increased significantly, resulting in additional rack clutter and greater potential for airflow blockages within the rack.

As a result, rack depths have increased to 42", and even to 48", to accommodate a deep server form factor with provisions in the rear (behind the servers) for power distribution, cable management and increased airflow.

Assuming that all variables relating to air blockage within the rack are constant, increased rack depth has a neutral effect on airflow. Although with heat containment, deeper racks (> 42") provide more space to exhaust air upward, as opposed to simply pushing it through the rear of the rack.

Although a rack width of 30" is common for networking applications with"side-to-side" cooled switches, the industry trend for server rack width is still 24". However, it is not uncommon today for 30" wide racks to be deployed for server applications. The wider server rack enables the end user to minimize airflow obstructions by locating power distribution and data cables further away from the hot air exhaust stream.

In most cases, room geometry will dictate rack height. The most common server rack height is approximately seven feet, which provides 42U of internal mounting available for rack-mounted equipment (1U = 1.75"). However, convergence of networking devices with servers is pushing the rack to grow to a height approaching eight feet (51U).

Additional rack-mounting space above 42U is typically populated with networking switches, routers and patch panels. Higher racks accommodate more rack-mounted IT devices. This results in even higher heat densities at higher elevations, where air stratification makes it difficult to ensure optimum inlet temperatures.

Rack as a Plenum

Today's hot and cold air containment solutions are highly dependent on a tight interface with the rack. Therefore, there is a need to change our thinking about the rack and its function in today's data center environment.

The rack should be thought of as a "plenum" in the airflow stream. Unlike a typical empty air duct plenum, the rack plenum is the critical space in which high-performance servers, storage and switches reside. To ensure sufficient IT device cooling, predictable rack level airflow management is necessary. In order to achieve this, all potential airflow openings should be controlled and managed. In addition to sealing unused U-space in the rack, there are at least five other rack-related areas that can directly affect airflow management and cooling performance, as well as improve energy efficiency.

Rack Hygiene Overview

The concept of Rack Hygiene is a term that encompasses the identification, analysis and repair of hot air leakage areas/infusion paths and cold air bypass routes within and around individual data center racks.

Rack Hygiene is a newly coined term used to describe the care in which the rack envelope is designed, controlled and maintained. The rack envelope consists of the entire volume of space from the floor to top of the rack itself; and perhaps, a measurement of rise above the floor to include empty space above the rack that would rise to the "heat deck".

Regardless of the rack's dimensions, it is incumbent upon the data center professional to employ solutions that provide an impenetrable barrier around the front plane of the rack in a front-to-back dominated airflow environment. The tighter the seal provided around the front of the rack – exclusive of a thorough blanking panel strategy – the closer one can come to achieving rack hygiene nirvana.

By addressing air leakage at several points around the front of the rack, one can reduce the supply path Delta T and the return path Delta T, resulting in higher energy efficiency. Therefore, air conditioning units will not need to work as hard to over-supply the demand as a means of eliminating data center hot spots and airflow abnormalities.

Two primary drivers of rack hygiene best practices are:

1. Hot spot prevention. This helps to maintain a constant inlet temperature and allows IT equipment to operate at optimal levels.
2. Matching the cooling supply and demand. This can save energy and eliminate wasteful recirculation and by-pass air streams which are part of the "chaotic-cooling method" (over-supply of cool air into the data center).

Ineffective rack airflow management at the rack and row level is a key contributor to aisle and room overheating. Rack Hygiene can solve this problem by approaching the rack as part of the airflow management system and setting benchmarked standards for leakage.

Five Airflow Fault Areas

Though the industry has learned the benefits of blanking panel best practices for the data center, this is only one airflow containment measure within the Rack Hygiene approach. There are as many as five additional rack-related areas requiring containment that are overlooked when one discusses rack-based airflow management faults. These fault areas can drive true performance gains in a front-to-back cooled world. The areas, known as the Five Airflow Fault Areas (Figure 2), include:

1. Under the rack (external to rack)
2. Left side of front-left 19" vertical mounting rail (internal to rack)
3. Right side of the front-right 19" vertical mounting rail (internal to rack)
4. Below the bottom rack-mount space (internal to rack)
5. Above the top rack-mount space (internal to rack)

Fault Area #2 & #3: Left and Right Side of Front 19" Vertical Rails

Because of customer demand for adjustable front rails and cable pass-thru capability, the areas to the left and right of the front rails on most 19" racks are potential leakage points. The space between the side of the vertical rail and the side of the rack frame or side panel is typically wide open. It is a potential leakage area into which hot air can penetrate or by which cold air can pass. This rack environment can severely compromise a robust blanking panel strategy.

Today's wider racks (i.e. 30") have an additional three inches on each side of the 19" rails to provide space for cooling side-to-side switches or space for managing a high volume of network cables. To get out to the side or up to the top, cables are passed through openings that are typically unsealed. These openings should be covered with a material that provides a seal around the cables to minimize air leakage.

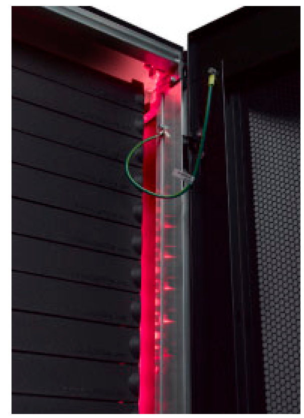

Fault Area #1: Under the Rack

The area under the rack to the floor deck can be difficult to manage because the height is a variable based on the size of rack levelers or casters and will vary from one rack manufacturer to another. This space can contain a substantial amount of uncontrolled air in an enterprise data center with multiple rows of server racks. Therefore, this is an area that can yield a large benefit, if sealed appropriately.

Typically, there is no solid panel under the rack due to the requirement for power and network connectivity. This is a potential leakage area because hot air can come across from below the rack, and cold air from perforated floor tiles can bypass the rack in this space.

Fault Areas #4 & #5: Above and Below the Vertical Rack-mount Space

Areas above the top U space and below the bottom U space are also regions of suspicious leakage. Typically, some amount of space exists in these areas and varies per rack manufacturer. However, it is not unusual for this space to equal that of a missing blanking panel.

Not only is this area susceptible to hot-air recirculation, but it is also more likely to allow bypass of the cool supply air supply from CRACs.

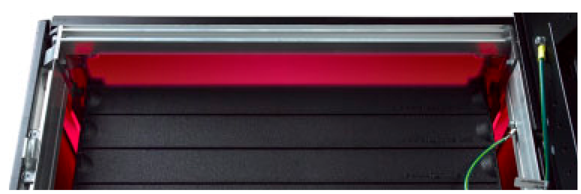

Measurement

Just as the efficiency of data centers is measured with PUE and DCiE, the efficiency of the rack should be measured to promote proper airflow management. The correct approach to rack hygiene and to reducing data center energy consumption is measurement. That is, establishing a baseline, and tracking performance for a given data center facility.

Benchmarking methodologies set performance goals for air leakage in the data center. By combining virtual models with measured test results, one can locate precisely where, in the data center, airflow management issues arise and take corrective action to improve airflow containment.

The benchmarking process will assign a graduated scale of performance to:

- Identify problems with energy profiles in the data centers
- Analyze data center energy performance for potential improvements
- Add high density servers and increase rack density in an energy efficient way
- Determine and select efficient data center cooling methods
- Predict design limitations of new and future facility expansion

In order to measure airflow leaks, data center design specification calls for server racks designed and supplied with a maximum rated amount of airflow (CFM) to service a full load rating.

The pressure within the racks should not exceed 0.001 inches/ H_20 at the front of the rack. Ideally, the IT equipment in a rack is exhausting air without experiencing any significant back pressure. Given this, the test pressure from the front to the back of the rack must be set not to exceed a maximum gap of 3% or less of the total surface area. Therefore, the pass/fail criteria would be established to require measured air leakage of up to 3% of the total supply air.

Testing

To test the racks and ensure they meet the required specification, the following tests are performed:

Measurement of the overall gaps in leakage area (they will be measured as a ratio to the overall surface area of the rack enclosure inlets.)

Measurement of whether the environment can maintain 3% or less leakage of air at 0.001" in H_20. (a fan will be used to pressurize sealed containment racks.)

Mapping and identification of leakage areas within the rack (a fog generator will be used to trace and detect areas of concern.)

Data center testing includes the Five Airflow Fault Areas:

1. Under the rack (external to rack)
2. Left side of front-left 19" vertical mounting rail (internal to rack
3. Right side of the front-right 19" vertical mounting rail (internal to rack)
4. Below the bottom rack-mount space (internal to rack)
5. Above the top rack-mount space (internal to rack)

Tests include those that determine:

- Withstand Pressure Range: To determine the envelope's useful operating range
- Leakage Level: To measure the amount of air that escapes in the Five Fault Areas when the U space is 100% blanked off

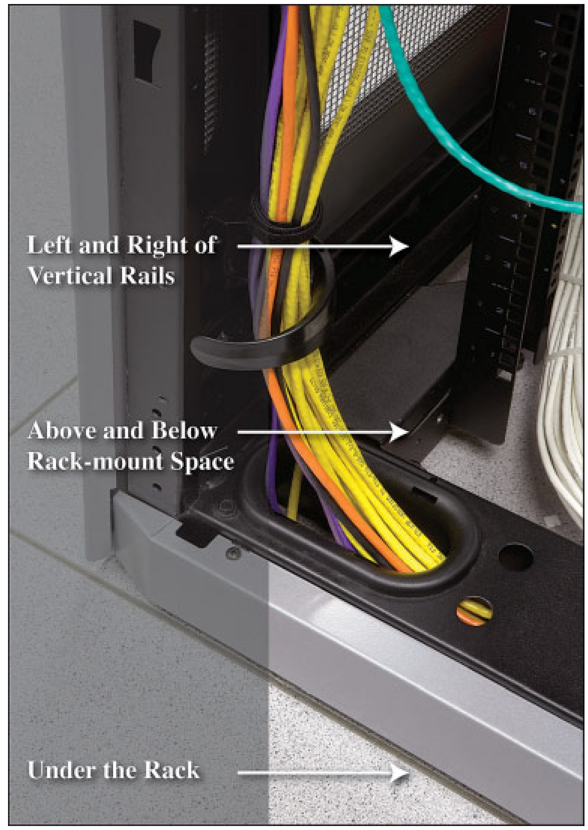

Figure 2. Rack fault areas.

The Goal is Zero Leakage

An experienced team, trained in containment strategies should analyze test results and provide a list of immediate energy efficient cost saving solutions.

Reporting includes:

- A summary overview of test results measurements
- A detailed analysis of each test including:
    - The facility, infrastructure, and baseline metrics for which the test were conducted
    - The apparatus used to conduct each test
    - The procedure used to conduct each test
    - How the results were collected and measured
- A spreadsheet indicating which racks have passed and failed each of the tests
- A summary illustrating any issues in which the specification was not met and where improvements would be required, equal that of a missing blanking panel

Conclusion

The evolution of the rack as a critical component of the engineered airflow system enables the data center to become more energy efficient, saves on costs and restores flexibility to the data center manager.

Unfortunately, the industry has relied on workplace intuition and creative problem solving for far too long. It has utilized everything from cardboard and duct tape to foam seal kits and other "weather-proofing" types of devices. However, the market is moving toward standardized solutions that are designed and integrated within the rack during the manufacturing process.

As data center containment at the rack and row level strengthens its foothold on retrofits and new construction, more data center managers will look for elements of Rack Hygiene to be included as standard feature sets, not premium options, of their future rack purchases.

The separation of hot and cold air dramatically increases the predictability of the data center's performance and enables:

- Efficient utilization of existing physical infrastructure and cooling capacity
- Active control and normalization of supply temperature, eliminating recirculation and stratification
- Doing more with less in a smaller data center footprint with increasing heat loads
- Elimination of "stranded" physical, electrical and mechanical capacity

A smart containment strategy begins with the rack, regardless of whether or not one currently is employing cold aisle or hot aisle containment. Improved Rack Hygiene, even in legacy chaos-cooled environments, is the first step toward mitigating the re-circulation and re-mixing of hot and cold air streams in the data center. Justifiably, it could account for up to 60% of the overall data center containment strategy.

About the Author

A technology industry veteran, Edward Eacueo, has 25 years of global experience leading sales, marketing, and engineering organizations in information technology, commercial electronics and military markets.

Eaton Corporation
Electrical Sector
1111 Superior Avenue
Cleveland, OH 44114
United States
877-ETN-CARE (877-386-2273)
Eaton.com

Publication No. 11-12/3-11

Powering Business Worldwide