A LAYMAN'S EXPLANATION OF THE ROLE OF IT RACKS IN COOLING YOUR DATA CENTER

By David Moss
Dell │ Data Center Infrastructure

dell.com/hiddendatacenter

Many people see an IT rack as a glorified shelving system. Structurally it is. It holds and separates IT systems, allows them to translate for service, potentially protects them from seismic events, enables high density stacking, and helps to make them look nice and orderly. Thinking past these physical or logistical attributes, the rack can make a big difference in how effective your room cooling strategy is. To understand the difference between a good rack implementation and a bad one, you can probably concentrate on the three most common questions I receive:

1. How many servers will this rack cool?
2. Are blanking panels really that important; how can I justify?
3. Isn't this serpentine bracket that holds all my cables causing overheating problems?

HOW MANY SERVERS WILL THIS RACK COOL?

Unless it is an actively cooled rack (i.e. one with a cooling coil, which looks like a radiator), the rack is a passive player in your thermal system. The amount of IT systems the rack can "cool" depends on the room's ability to deliver adequate temperature air to the rack.

Each IT system consumes a specific amount of air. This air enters the system through the front, picks up heat while it passes through the system, exits the system through the back, and eventually exhausts from the rack. A good system vendor can tell you how fast a particular piece of IT equipment consumes air (in the US, this is generally expressed in Cubic Feet per Minute or CFM.) Ideally, the facility delivers cool air to the systems in volumetric amounts that match system CFM ratings. If the facility is incapable of supplying sufficient volumetric amounts of "cool" air, the IT equipment will draw in air from places you'd rather it not. This typically ends up being used exhaust (i.e. hot air) from somewhere in the room and generally occurs in systems that are furthest away from cool air vents.

So, it is the room's inability to deliver sufficient quantities of air that results in some systems receiving hot (and often) out of spec air temperatures; and it's the room's ability that establishes how much IT equipment can be cooled, not the rack's ability.

WHY SHOULD I USE BLANKING PANELS?

A decade ago, IT rooms were cooled much like your home or office. Even though there may have been a greater concentration of air conditioning, the strategy was pretty simple – "cool the room". Air was ducted into the room but with no particular strategy as to how it entered or left IT equipment. Back then, empty rack spaces really didn't need to be closed off by blanking panels. As IT spaces evolved, it became apparent there was a better way to put more systems in the same space. If you could concentrate the air conditioning and duct it close to the inlets of the IT equipment, you could cool systems with less air conditioning (or put more systems in the same room). Thus, the hot/cold aisle strategy was born.

Hot/cold aisle orientation creates separation between hot air and cold air in a data center. Intentional containment systems take this separation further. In both cases, the need to prohibit re-circulation of hot air within the rack becomes much more critical. With today's high performance IT systems, exhaust temperatures are generally significantly higher than with older IT systems and much higher than the inlet specification for most IT equipment. An open "U" gap is a "short circuit" that allows re-circulation to affect the inlet temperature of adjacent equipment. A simple CFD model can explain the impact of an open "U" gap.

This model was created for purposes of cost justification for installing a blanking panel.

- A rack is loaded with forty current day 1U servers of typical configuration. 68 degree air is supplied to the entire front surface of the rack (left side) to simulate a room that is providing adequate cooling.
- If one of the servers is replaced by a gap, and the gap is not filled with a blanking panel, then the inlet temperature of several adjacent servers is affected. The two servers surrounding the gap operate at an average inlet temperature of 95 (the upper allowable limit for these servers). The next two adjacent servers also experience elevated temperatures in the 80s.
- To compensate for the increased temperatures, the increase in fan power for these four servers is about 185 watts. Although this is not an enormous increase (just over 1%) for a rack using nearly 13 kW, the increased power adds up over a year. At $0.10 per kW-hr and a PUE of 2.0, an additional 185 watts results in $324 of added electrical consumption.

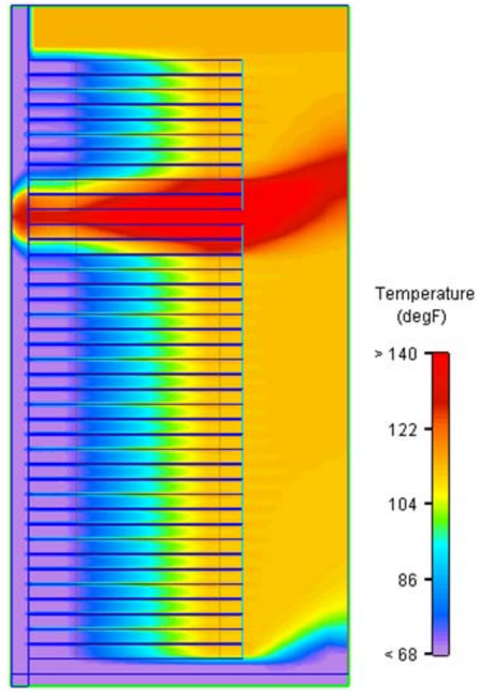

So, is a $5-$10 blanking panel, which can prevent such re-circulation, worth it? Probably!

ISN'T THIS SERPENTINE BRACKET THAT HOLDS ALL MY CABLES CAUSING OVERHEATING PROBLEMS?

To better understand the answer to this question, you first need to have a better understanding of the physics of how a rack works. An ideal rack makes a good seal between its front flanges and the front IT brackets. This allows the cavity between the IT systems and the front perforated door to become negatively pressured. An obvious detriment to an effective seal would be any gap between the front and the rear of the rack (i.e. between systems, laterally in the rack, or missing equipment without blanking panels). The negative pressure front cavity is what enables fresh, cool air to be drawn into the front perforations. Likewise in the rear of the cabinet, it is the positive pressure in the rear of the cabinet that forces air through the rear door to an area of lesser pressure (the hot aisle). Again, gaps between systems or around the rack flanges lower the rack's ability to push air through the rear door. CFD will be used to illustrate what the cable arm blockage does and how the air finds its way out of the rack.

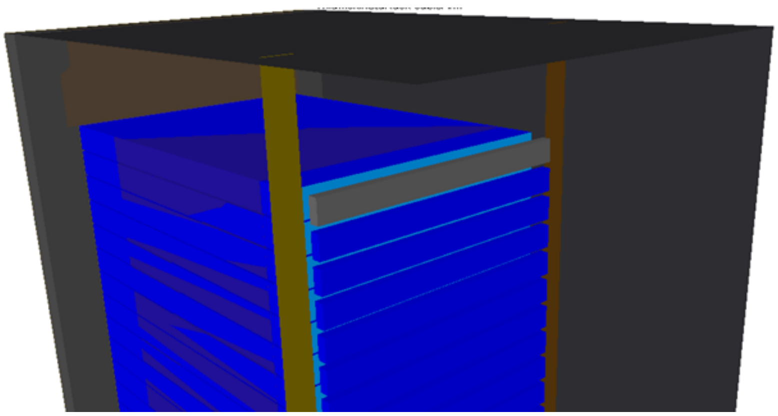

The cable arm was modeled by adding a solid blockage behind each system. There are small gaps between the arms and between the arm and the rack flanges. There is about 2 inches between the outside of the flange and the inside of the rack side panels that air can travel around toward the rear of the rack.

The figures below depict the pressure at different portions in the rack. Relative to a zero pressure outside the rack, the purple areas show where the pressure is negative. Air will migrate from outside the rack into the purple area. As the air exits the server, the pressure becomes positive relative to the pressure outside of the rack. A high pressure (the red area) is seen between the server and the cable arm blockage. This high pressure is what slightly slows the flow through the server. Assuming there is an adequate path to the rear (laterally and between cable arms), pressure still builds up in the rear of the rack similar to the levels seen without the cable blockage. Air is still exhausted through the door because of the pressure differential across it. There is very little difference between the pressure profiles of the rack with and without the cable arm.

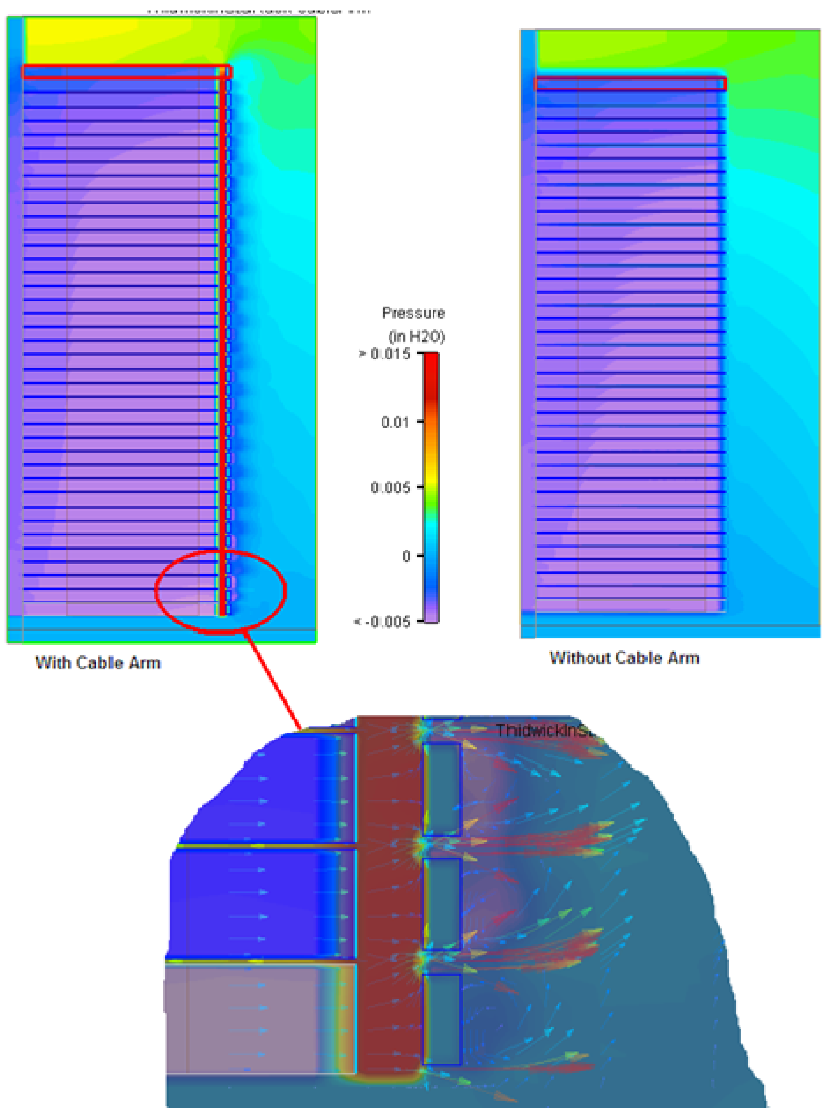

We feel that the cable arm resistance can be considered negligible unless there is no alternate path for the air to travel rearward. There is an assumption that some of the air will pass through gaps between the arms but that most of it will travel laterally around any rack flanges and make its way to the rear cavity. If the cable arm lines up with rack flanges and the adjacent zero-U space (between the IT equipment and the rack side panel) is filled, this would be a scenario where significant blockage might occur. The zero-U space is often used to mount PDU strips. If the PDU, rack flanges, and cable arm line up (in the same plane), and it is significantly blocked up and down the height of the rack, this could impede air enough to be of significance.

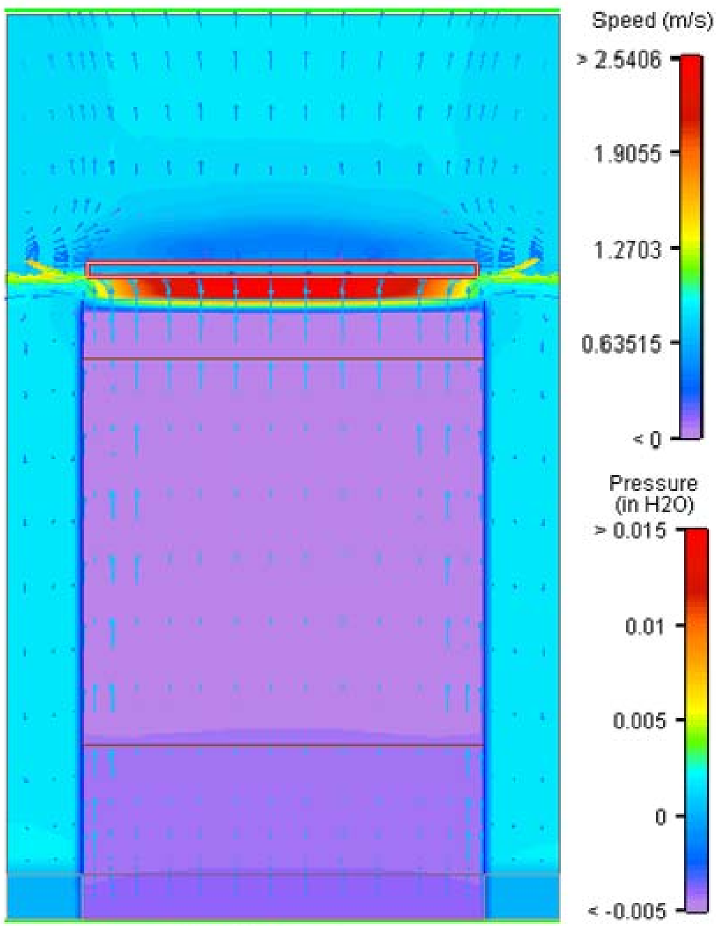

View from Top

In a recent Dell study[1], the flow resistance was measured through door perforations and around loaded cable arms. The study found that the resistance attributed to a cable arm is about the same as the resistance of a front and rear perforated door combination.

1. "Rack Infrastructure Effects on Thermal Performance of a Server", KC Coxe, 2009

In this graph, you will see how the resistance (impedance) increases (exponentially) as more air is passed through any of the restrictions considered. The red line represents a typical server or fan response. As less backpressure (impedance) is applied to it, it produces a greater flow rate. The intersection of it with any of the impedance curves would represent the effect of the impedance on the flow rate for the server. In this case, if you close one door there would be about 34.5 CFM passing through the system. If you closed both doors and added a cable arm, the system flow rate would drop to about 34 CFM. Two closed doors and the cable management arm would cause a decrease in the server flow rate to about 33 CFM.

IT vendors take into account typical external resistances that might be encountered by equipment and make sure that net flow rates still produce acceptable temperatures.

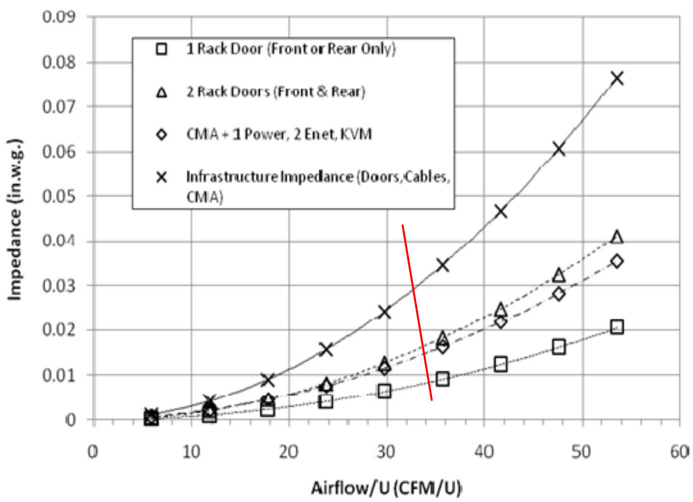

Think of it as nothing more than an added pinch point. Air traveling through the rack squeezes through many restrictions between the front and the rear. It squeezes through door perforations, the front bezel vents (if present), the front chassis vents, very restrictive heat sinks, rear chassis vents, and eventually the rear door perforations. It is the sum of these restrictions, including a cable arm, which establishes the resulting flow rate. Even something as apparently restrictive as a full cable management arm is but an incremental resistance when you consider flow through the entire rack. It's kind of like a kink in your garden hose. After you straighten it out, you might still see a noticeable indention. It probably is slowing the flow down a little bit. But other factors probably play a heavier role in deciding what the ultimate flow rate is- other factors like hose length, hose diameter (I'll never let my wife buy another ½" hose), what kind of sprinkler, etc. That kink is just one resistance of many contributing to the resulting flow.

In summary, for anyone concerned about the thermal effect of a cable arm, they should be equally concerned about closing their front and rear doors. Our testing concludes there is a slight internal component temperature rise due to all of the rack resistances: front door, rear door, and cable arm. This added resistance slows the air down slightly, 6% (half due to the cable arm), but this is expected and should be of no concern relative to the thermal or compute performance of the IT equipment. The decision to use a cable arm should be made on the basis of serviceability. Is it needed to provide a service loop when extending the equipment for component replacement? If you are a "rip and replace" shop, you probably don't need one; why pay for it? But if you do service in the rack, you should not concern yourself over the minimal thermal impact associated with a cable arm. Results similar to the Dell testing may also be seen in tests published from HP (Rubenstein, B (2008). "Cable Management Arm Airflow Impedance Study").

For more information, please consider the following links:

Rack Infrastructure Effects on Thermal Performance of a Server:
http://content.dell.com/us/en/corp/d/business~solutions~whitepapers~en/Documents~dci-rack-infrastructure-effects-on-thermal-performance-of-a-server.pdf.aspx