Data Center Air Management Metrics-Practical Approach

by Munther Salim and Robert Tozer

Perhaps the only thing growing faster than energy use in data centers is new metrics and acronyms to measure and control them. One way to consider the problem of data center energy use is to view the data center as a giant air pump. In a perfect world, every CFM of air produced by the CRAC units would find its way to the server inlets.

Unfortunately the world is not perfect. This makes air-management metrics a powerful tool for improving the energy efficiency of data centers. By quantifying airflow performance, facility executives can measure and quantify the effectiveness of their data center cooling systems or changes that they make to their cooling systems to increase air cooling efficiency. This approach applies to both traditional raised floor designs and non-raised floor air-conditioning designs for data centers

This article presents means to quantify air flow performance of both traditional raised floor designs and non-raised floor air-conditioning designs for data centers. Metrics are developed that will readily permit owners, engineers and operators to measure and quantify the effectiveness of their data center cooling systems or changes that they make to their cooling systems to increase air cooling efficiency.

This article proposes novel data center air-management metrics that apart from re-circulation, also address negative pressure flow, bypass flow and balance of server and CRAC unit flows. The metrics are practical and simple to apply and provide a specific characteristic for each data hall.

These metrics establish how well the cooled air from the CRAC units reaches the server inlets. Ideally, all of the air from the CRAC units would go to the servers. With these metrics, facility executives can create a picture showing how far from the ideal a data center is, and use that picture to improve air management.

For these metrics, the main parameters are air mass flow rates and temperatures, which are used to express the quantity and quality of energy supplied to the servers. These air management metrics are based solely on the consideration of sensible loads. For this conceptual model, the following flow rates are considered.

CRAC flow: Total air-flow rate produced by all operating CRAC units in the data center is normally greater than that required by the servers, due to part load operation of IT and communications equipment, and redundant CRAC units that are kept operational.

Bypass airflow: Air that leaves the floor grilles and returns directly to the CRAC unit without cooling servers.

Re-circulation airflow: Air discharged from servers that returns and mixes with chilled air entering the servers.

These are indicated in Figures 1-4:

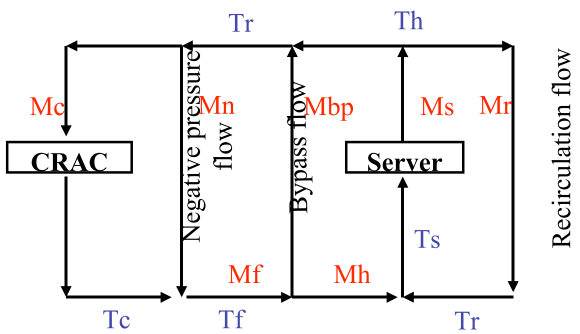

Figure 1: Data center air streams

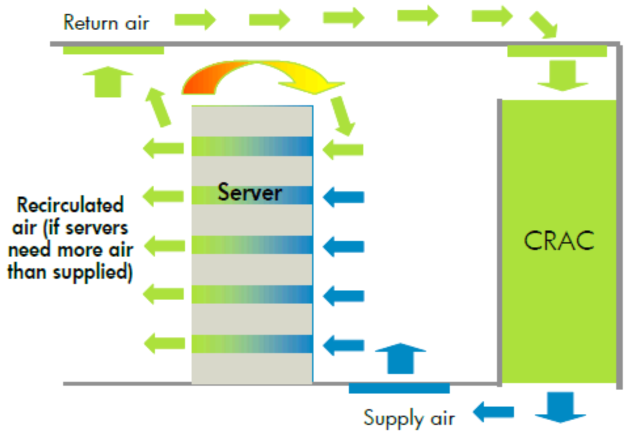

Figure 2: Recirculation air flow

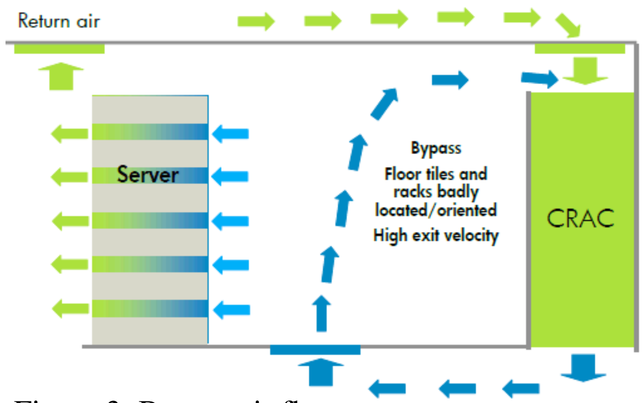

Figure 3: Bypass air flow

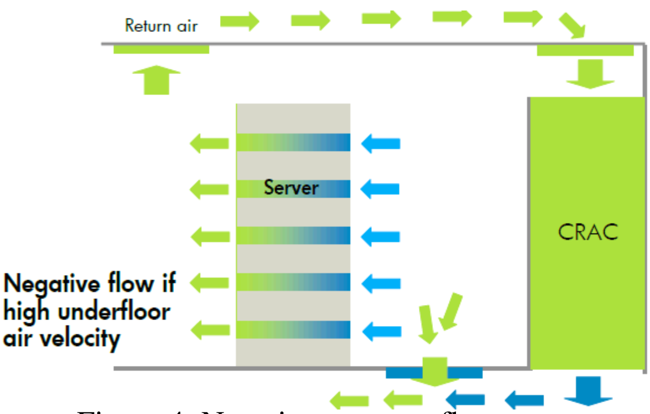

Figure 4: Negative pressure flow

Temperatures as found in typical legacy data centers are used:

- T_r = 21°C, return air temperature to CRAC (at CRAC), normally the CRAC return air set point
- T_c = 14°C, discharge air temperature from CRAC (at CRAC)
- T_f = 14.1°C, floor void temperature (after room air is drawn in), very close to T_c
- T_s = 21°C, server inlet air temperature (mixture of grille and re-circulation air)
- T_h = 28°C, server outlet air temperature (before cold air mixes with it)

Using these mass flow rates and temperatures, the following mass (3 nodes) and heat balance equations can be written.

Mass equations:

$$m_f = m_n + m_c$$
$$m_f = m_b + m_h$$
$$m_s = m_h + m_r$$

Heat equations:

$$Q_c = m_c c_p (T_r - T_c)$$
$$Q_f = m_f c_p (T_r - T_f)$$
$$Q_h = m_h c_p (T_h - T_f)$$
$$Q_s = m_s c_p (T_h - T_s)$$

From these equations, the following data center air management metrics are derived.

Negative pressure ratio:

$$NP = \frac{m_n}{m_c} = \frac{T_f - T_c}{T_r - T_f}$$

Bypass ratio:

$$BP = \frac{m_b}{m_f} = \frac{T_h - T_r}{T_h - T_f}$$

Re-circulation ratio:

$$R = \frac{m_r}{m_s} = \frac{T_s - T_f}{T_h - T_f}$$

Balance ratio:

$$BAL = \frac{m_c}{m_s} = \frac{T_h - T_s}{T_r - T_c}$$

The relationship between Negative Pressure, Bypass Pressure, Recirculation and Balance, Ratios

Using the concept of equivalent balance between the floor flow (CRAC flow plus negative pressure flow) to server flow, the following expression is derived:

$$BAL(1 + NP) = \frac{m_f}{m_s}$$

If these mass flow rates are expressed in terms of bypass and re-circulation mass flow rates, and with further derivation gives:

$$BAL(1 + NP) = \frac{1 - R}{1 - BP}$$

This relationship is represented in Figure 5.

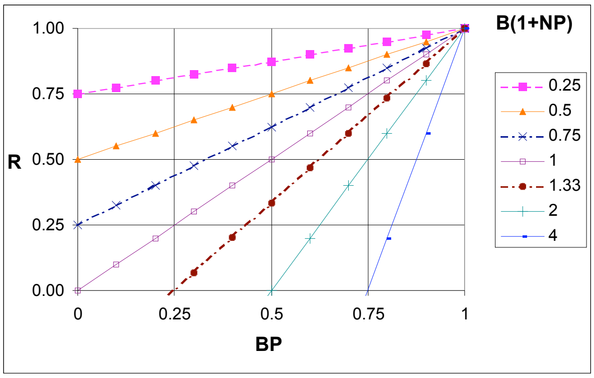

Figure 5: Data center air management metrics

This graph displays common values for recirculated and bypass airflow ratios for different air-management strategies.

The lines on the graph represent the balance between air supplied by the CRACs and air required by the servers. (The numbers in the box to the right of the graph are the values of the lines.) The pink line at the top of the graph corresponds to a balance of 0.25, which means servers are not getting enough cold air, while the blue line on the far right corresponds to a value of 4, which means the CRACs are producing four times more air flow than the servers require. The purple line in the middle means sufficient airflow. For values above the purple line, there is not enough air going to the servers; while below the purple line, the opposite problem exists — too much airflow.

Balance shows the ratio of air from the CRAC units that reaches the server inlets. The ideal balance is 1, showing that all the air from the CRAC units is going to the server inlets.

The scale for these ratios is zero to 1. For bypass and recirculation, the ideal result would be zero — that is, no bypass air and no recirculation.

On the other hand, if the bypass ratio is 1 (BP=1), all of the CRAC cooling air would be bypassed, with no air available for servers. And similarly, if recirculation is 1 (R=1), all server air would be re-circulated, with no cooling air entering from the CRAC units.

Practical Applications

Even if it is difficult to precisely measure NP, BP, R and BAL, there is always a real specific value that is characteristic of every data center.

Negative pressure flow is normally very low and difficult to measure, however, it could be effectively modelled and reported using CFD (computational fluid dynamics).

In practical terms, the CRAC system should be controlled based on supply air and not on return air. Unfortunately, few data centers are operated that way. Controlling on supply air improves the uniformity in server intake temperatures, avoids unplanned latent cooling, boosts the uniformity of chilled water openings on the chiller water CRAC units and reduces fighting among controls.

While retrieving exact measurements is difficult, the authors obtained representative indices for different types of air management solutions. These are indicated in Figure 6.

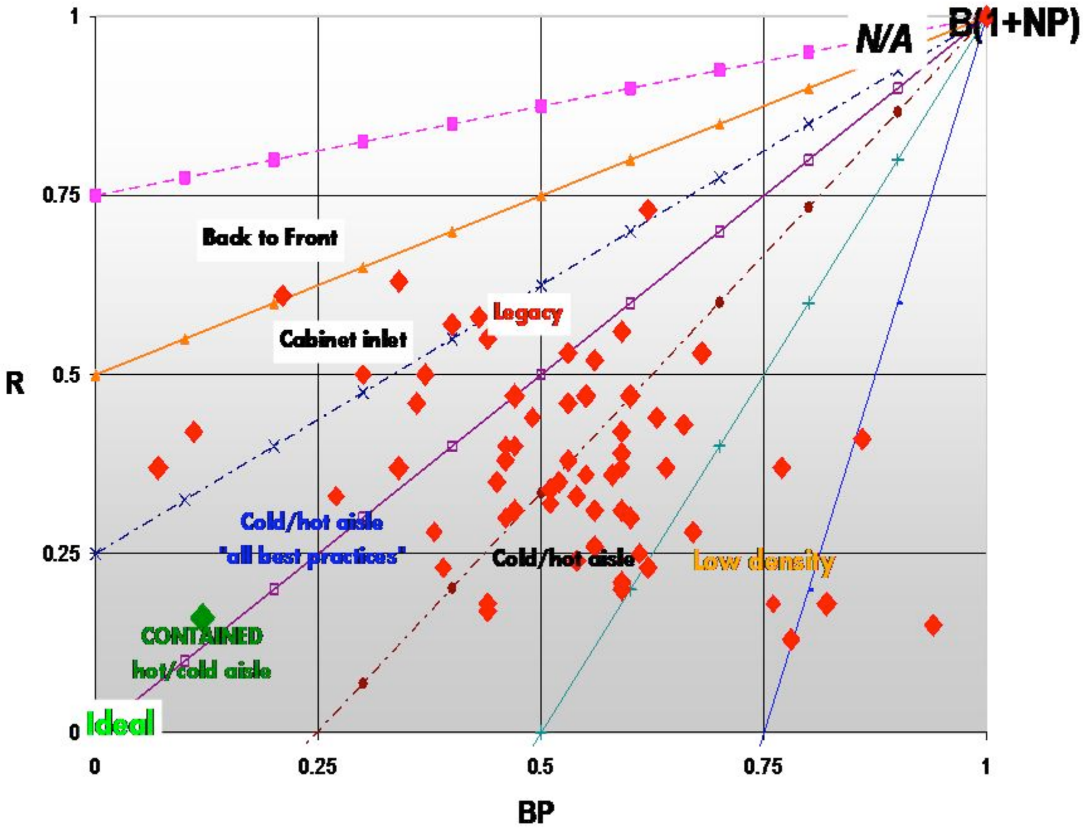

Figure 6: Air management solutions

The red diamonds show measurements from actual data centers.

• Legacy data centers, particularly if not well-managed, tend to have high levels of both bypass and re-circulation flow rates, which are increased when subject to higher density loads.

• Cabinet inlet refers to closed cabinets with air entry from below and with fans on top of the cabinet to extract heat. Generally there are no blanking plates within the cabinets, and inlet and exhaust air streams are not segregated, producing high levels of re-circulation. Bypass air is better managed unless there are a large number of floor grilles, which would bypass air back to the CRAC units.

• Back-to-front arrangements tend to have low bypass but high levels of re-circulation from the back of the cabinets to the front of the cabinets in the same aisle.

• Cold-aisle/hot-aisle arrangements, if properly designed and operated, can address the main problems of re-circulation. However, particularly at low loads, there will be high levels of bypass air, unless some type of variable air-flow control is fitted.

• Cold-aisle/hot-aisle schemes with all best design practices — such as CRAH units in galleries with high bay ceiling (or return ceiling plenum with ducted CRAHs), blanking panels, cable brushes, foam pillows, and continuous rows — have shown reduced recirculation and bypass levels. However, a big opportunity still exists.

• For high density loads, there are several designs whose basic intent is to contain and separate the cold air from the heated return air on the data floor: hot-aisle containment; cold-aisle containment; contained rack supply, room return; room supply, contained rack return; contained rack supply, contained rack return. The separation can be achieved with hard surfaces “glass” or strip curtains for example. Contained cold/hot aisle is a relatively new approach and performance data indicated very low levels of bypass and recirculation.

Improving Performance

Ideally, the data center would have zero bypass and recirculation air flow, and there would also be a balance between the CRAC and server air requirements (BP=0, R=0, NP=0 and BAL=1). But that is far from the case in reality. From surveys and assessments carried out by the authors, on average, half the air flow rate produced by the CRAC units bypasses IT equipment and returns to the CRAC units without doing any useful cooling. And approximately half the air intake to the servers is exhaust air from servers.

Using the discussed air-management metrics, detailed calculations were made for a mid-range data center with bypass and re-circulation factors of 0.5. Eliminating airflow bypass can reduce CRAC fan energy use by 60kW. By eliminating re-circulation, compressor energy savings of 74kW are possible, due to the evaporator operating at higher temperatures. For a 10,000 square foot data center, this can represent a saving of $90,000 annually. These estimates do not include free cooling, which has the potential to increase savings by an order of magnitude.

How can facility executives improve the air management of their data centers? From the practical applications the following is a list of measures that reduce the issues of poor air management. They are listed in descending order of impact and feasibility:

<u>Bypass Air:</u>

• Seal air gaps in the raised floor, particularly cable cut-outs within cabinets at the discharge side of servers.

• Relocate floor air grills so that they supply the inlet of servers — that is, remove floor air grills from hot aisles and other areas where cooling isn't required.

• Ensure air velocities from floor grills are not too high, which can cause air to overshoot the top of the cabinet.

• CRAC units turned off are a source of air bypass; consider air isolation with dampers.

<u>Recirculation Air</u>:

• Fit blanking plates in cabinets where servers are not installed.

• Close gaps between cabinets where warm air can make its way to the server inlet.

• Ensure as much as possible the adequate supply of cold air to the server inlets.

• Ensure air velocities and flows from floor grills are sufficient to reach the top of cabinet servers.

• Review air return path to CRACs and consider return air plenum or ductwork if necessary.

• Consider means to physically isolate the supply (cold) and return (hot) air streams.

• Remove obstructions under the floor that restrict cold air to supply the server inlets.

The effectiveness of those measures can be seen from work done at a 5,000-square-foot data center. During an energy assessment of the facility, measurements were taken from more than 20 percent of the servers present and all of the CRAC units. The general data center floor layout is back-to-front, but some server discharges and intakes are on the same side. The load density was 27 watts per square foot. Recirculation, bypass, and balance were calculated to be 0.61, 0.21, and 0.50, respectively. The load density was 27 W/ft^2. Recall that the ideal for recirculation and bypass is zero, and for balance it is 1. Clearly there was room for improvement.

Interestingly, when these points were plotted on Figure 4 above, the point fell into the area marked by back to front. The balance ratio of 0.5 indicates that the servers are not getting enough air. That was attributed to the fairly low height (18 inches) of the raised floor with many obstructions. The obstructions resulted in low air velocities from the floor grills, so that cool air didn't reach the servers in the top of the cabinets. Another problem is that many server intakes and discharges are located in the same aisle.

To reduce recirculation and increase balance, the assessment team recommended better management of underfloor obstructions; more efficient layout such as cold/hot aisle; retrofitting the data center with a return air plenum; blanking-off all the unused server locations; fitting blanking panels between cabinets; and fitting brush seals at all cable cutouts. Based on measurements taken from the data center at the same locations after implementing most of the recommendations, new air management metrics were 0.34 for recirculation, 0.18 for bypass and 0.82 for balance, a substantial improvement. (A negative pressure ratio of 0 is assumed.) The facility manager is looking at installing a return air plenum to reduce the recirculation and bypass flow ratios even more.

Lessons Learned

The following lessons can serve as guidelines for such analyses.

• For all temperature measurements use a simple, reliable thermometer, not laser guns since they measure surface temperatures and not air "ambient" temperature

• Measure and average the actual air outlet temperature (supply from CRAC) and air inlet temperature (return to CRAC) in a couple of locations across the duct and not only at a single point. In some cases, stratification may lead to different temperature values.

• The assessment team should sample two to three times as many servers as there are CRAC units, measuring the server air inlet and outlet temperatures.

• The server air inlet temperature should be taken with a thermometer as close as possible to server grill inlet (not at the rack or cabinet inlet) without touching any surfaces. Similarly, server outlet temperature should be taken as close as possible to the server without touching any surfaces.

• Server inlet and outlet temperatures should be taken at the same level for the same server. If temperature is measured at the server inlet located in the middle of the rack, the server outlet temperature should be taken for the same server at the same level.

• Readings should be taken from servers in a mix of locations: beginning of aisle, end of aisle, middle of aisle, bottom of cabinets, middle of cabinets, and top of cabinets.

• Readings should be taken from several aisles and not just one or two aisles.

• Measurements should be taken from several hot spots, if hot spots exist.

• It is acceptable to use average temperatures rather than weighting them with mass flow rate/size.

Many data center operators take a relaxed attitude toward bypass air, which is produced by the CRACs and returns directly without cooling servers. Bypass air causes a gap between the temperature of air produced by the CRAC units and the temperature of the air at the inlet to the server, although in theory those two should be the same. CRAC units are designed to produce lower temperatures for the sole reason of offsetting the re-circulation of air from the outlet of the server to its inlet. Air bypass also starves servers of air, which increases recirculation.

Once negative pressure flow (NP), bypass flow and re-circulation flow are addressed, the balance of CRAC and server air (BAL) becomes a necessary issue to address. This will require some method of automatic balance of air to load.

Ideal future air cooled solutions will require variable CRAC air flow rate as well as segregation between cold air to servers and discharge warm air from servers.

Better air management within data centers — mainly minimizing bypass and re-circulation airflow — has the potential to improve temperature control for more reliable server cooling. It can also save energy. What's more, improved air management is essential to maximizing data center free cooling opportunities.

Munther Salim, PhD, is principal, energy efficiency and sustainability in the Chicago office of HP Critical Facilities Services - EYP MCF. Robert Tozer, PhD, is technical director - HP EMEA.