



Donald L. Beaty

## Part 5: Case Studies

# Changing Landscape of Data Centers

BY DONALD L. BEATY, P.E., FELLOW ASHRAE; DAVID QUIRK, P.E., MEMBER ASHRAE; JEFF JAWORSKI

Addressing disruptive change in the data center can be a complicated process and challenging to predict. As rack density trends upwards and cooling becomes a greater challenge, there is more potential for both overcooling and undercooling of IT equipment.

Simultaneously, IT workloads are trending towards being more on demand, and thus less predictable. The speed at which loads can be increased or shed, in order to meet demand, may be the fastest that it has ever been, and it is happening more frequently.

It is therefore important to understand the changes that may take place within the data center on an hourly, daily, or monthly basis. It is equally important to understand (and if necessary, tweak) how the data center infrastructure should perform.

The previous columns in this series have discussed IT equipment, its thermal design, its interactions with the facility, and current and future sources of change within the IT space. This column will discuss how to identify, analyze, and predict change by applying the practical methodology of a case study to existing data centers.

## High Density Data Centers—Case Studies & Best Practices

The case study mentioned here can be found in the Book 7 of the ASHRAE Datacom Series, entitled *High Density Data Centers – Case Studies and Best Practices*.<sup>1</sup> Published nearly 10 years ago, this book takes an empirical approach to understanding the state of data center cooling in order to provide guidance for the high-density environments to come.

Today, we see high-density data centers around every corner, the majority of which use the very same cooling

approaches that were investigated in this book. There have been many improvements across the past decade, but the fundamentals remain the same. Further, the techniques employed throughout the book can still be applied today in order to study nearly any data center.

## The Case Study Process

In a typical case study of a data center there are four steps:

- Defining the study;
- Designing the study;
- Data collection; and
- Analysis.

The combination of these steps can provide immense value; skip a step, and get significantly less value.

The first step of a case study is to define the questions that are going to be asked and hopefully answered. This is probably the most frequently skipped step, but possibly the most important. Understanding the questions to be studied is necessary to ensure that the subsequent steps are done completely and without bias.

The range of potential questions for investigation is nearly limitless. One might investigate the causes of hot spots in a small section of the data center, or attempt to determine why the efficiency of one data center is

Donald L. Beaty, P.E., is president, David Quirk, P.E., is vice president, and Jeff Jaworski is an engineer at DLB Associates Consulting Engineers, in Eatontown, N.J.

exceeding the efficiency of another. It could be to investigate the cause of systemic equipment failure, or to simply understand whether the current data center can handle a refresh with higher density IT equipment.


The second step is to design the study. Armed with the clear goal in mind, it should be determined what data, measurements and observations will be necessary. This may include some form of spatial, power and cooling measurement, data regarding the installed IT and facility equipment, or observations of the qualitative conditions within the data center.

At this point, the scope of the study should be considered. If one were armed with all the data imaginable, there would be few questions that could not be answered. However, gathering data is no instantaneous task, and can easily get out of hand given the scale of many data centers. The sampling size or frequency may be scaled back, yet still offer sufficient data, if scaled back correctly.

Fortunately, as data centers have evolved, the number of installed sensors has grown dramatically, and typically offers up data at a scale that was previously unheard of. The wealth of available data may require modern analysis tools in order to process, but the speed and frequency at which this data can be collected can be invaluable... or it can cause analysis paralysis if managed incorrectly!

The third step is to proceed with the data collection as determined in Step 2. While this sounds simple enough, the organized, methodical and error-free collection of data should not be taken too lightly. The simplest mistake could easily contaminate data and lead to drawing the incorrect conclusions. For example, if airflow were to be measured throughout a cold aisle, measuring at variable heights or distances from the front of the rack could introduce dramatic variations in the data. *Figure 1* demonstrates some reference diagrams, which are helpful in the field to ensure measurements are done consistently and data is recorded properly

An enabler of consistency and thorough



documentation is the proper organization and recording of data. This is especially important for gathering data in the field, as it is easy to get “in the weeds” and make a mistake. Having predesigned input forms, spreadsheets, and graphical tools (annotated plans, key plans, floor grids, etc.) greatly increase the chance of accurate data, not only at the time of collection, but potentially days or weeks later when revisiting the gathered data.

Another easily overlooked consideration is a clear and thought out nomenclature. While it may seem simple, consider that even the basic orientation of left and right can reverse when observing both the front and back of a rack. Sloppy nomenclature could easily muddy field observations that could be important. Clear nomenclature also makes it easier to correlate measurements to drawings, making it easier to pinpoint where measurements were made.

The final step is analysis. Here, the effort organized into the first three steps culminates, into some well qualified answers. It would be difficult to cover the depth and breadth of analysis techniques that exist, and so we leave it up to the reader to research and apply the tools at their disposal.

However, analysis does not need to be complicated. Sometimes the more straightforward path the best answer. For example, if wanting to understand the uniformity of air distribution throughout an area of perforated floor tiles, little more than observation of the measurements may be necessary to find areas of

high and low flow. Add in knowledge of the layout within that space and it might be enough to form a reasonable hypothesis for the uniformity or lack thereof.

On the other hand, the analysis may be complicated involving a sensitive process with many variables and many field measurements combined with the data from thousands of sensors. Ideally, a well thought out formulation of the study will offer a clear path for analysis, rather than chaotically wading through mountains of data.

### Case Study Review

Understanding the case study as a process is an extremely valuable tool for investigation, not only within a single data center but potentially across many data centers. To provide a more practical understanding, the following will review an actual case study from the *High Density Data Centers* book.

#### Defining the Question

This particular case study deviates from the norm in the sense that it was for a book, and toward that end, the question being asked was simply about understanding the overall thermal profile of the data center. In a broader sense, the question could also have been to understand which of all of the data centers being compared was most efficient, and why.

#### Designing the Study

The design of the case study began by understanding the general nature of the data center in order to identify the measurements and data that would be collected. This case study was conducted at the National Center for Environmental Prediction. The data center consisted of approximately 6000 ft<sup>2</sup> (550 m<sup>2</sup>) of raised floor space, cooled by perimeter computer room air conditioner (CRAC) units.

Located on the raised floor were 51 racks of servers in five rows, with four of them arranged in two hot aisles. The majority of racks were low density, with a few racks with four times greater density. Perimeter power distribution units supplied power to the racks from beneath the raised floor.

Figure 2 shows the general floorplan of the data center. To provide an efficient system for capturing data, the raised floor tiles were used to create

a coordinate system, using letters along one axis and numbers along the other. In this way, each raised floor tile within the data center has a simple and unique identifier that helps to correlate a measurement to its location. The beauty of this simple technique is that the grid is physically visible while in the data center.

Physical data to be collected or verified included:

- The height of the raised floor;
- The level of congestion beneath the floor;
- The locations and orientation of equipment;
- The type and location of the perforated tiles; and
- The flow through them.

Power measurements of as much of the equipment as possible, including lighting, would aid in understanding the total heat load throughout the space. Air temperatures were measured at the CRAC supply, the server inlets, and the CRAC return to understand the temperature throughout the space.

#### Collecting Data

Armed with a plan for what was to be measured and how it should be measured greatly increased the chances that data collection would be smooth and free from error. This includes having the proper tools on hand for the necessary measurements and logging.

For airflow measurements, a properly calibrated velometer was used, which measures the flow of air through its hood. The velometer's hood was exactly the size of the perforated tiles, making it an efficient tool for the job.

For power and temperature measurements, high quality hand-held tools were used including a voltmeter, clamp-on current meter, and digital thermometer. Additional power data was gathered directly from servers using a laptop and specialized software.

These tools were the right ones for the job given the size of the data center and available manpower. However, methodology is just as important. Standards were observed where applicable, such as the ASHRAE guideline for measuring inlet air temperature 2 in. (50 mm) from the inlet.<sup>2</sup>

We leave it to the reader to examine this case study in full for a thorough description of the applied measurement techniques and methodologies. Suffice it to say that it is highly detailed and methodical. Further, many techniques have advanced since this case study and it is

recommended the reader carefully examine new approaches, tools, etc.

### Analysis

A great takeaway from this case study was that the analysis did not produce some of the expected results. For example, expected correlations between the airflow and air temperature rise at the racks were not found. However, other data yielded unexpected insight that airflow through the cable cut-outs was amounting to one-third of the total airflow, no doubt a large source of cooling inefficiency.


An important step in the analysis was validation of the results through an energy balance. An overall data center flow rate was calculated using the total dissipated power and differential temperature across the CRAC, and compared to the measured CRAC airflow. The compared values were within 5% of each other, providing a reasonable confidence factor in the validity of the measurement and analysis.

Undoubtedly, this case study revealed some detailed results about the data center's thermal profile that would not have been the same as the day the data center came online. The obvious variable like perforated floor locations, rack locations, and servers within the racks affect the performance of the data center, but so do other factors that are often overlooked. In this case, a lower than expected flow rate through the CRAC was attributed, in part, to dirty air filters, indicating that perhaps regular maintenance was not being observed.

### Conclusion

Change, in today's data center environment, is unavoidable. Starting from the day that a data center is commissioned, changes will occur that create differences between its intended and actual operation. Not only does the facility equipment change over time, but

FIGURE 2 High level data center floorplan with floor tile coordinate system.



the IT loads are changing more rapidly than ever.

The case study process is a helpful way to take back control of the changes happening to the data center. This does not mean that suddenly you can control the change, but rather understand how those changes are propagating throughout the data center and adjust accordingly.

Whatever the case, understanding the causes of change is as important as knowing how to mitigate it. By fully understanding the environment, the course of the data center's evolution can be altered in subtle or dramatic ways that can increase its performance, stability, or reliability.

### References

1. *High Density Data Centers; Case Studies and Best Practices*, ASHRAE 2008
2. *Thermal Guidelines for Data Processing Environments, Second Edition*, ASHRAE 2009. ■