IMPROVING DATA CENTER EFFICIENCY AND RELATED IMPACTS

ELECTRIC POWER CONTINUES TO BE CHALLENGING FOR THE DATA CENTER INDUSTRY

The energy required to operate the data center is still a major challenge. Large Data centers in particular, have invested heavily to optimize power usage, in addition to utilizing alternative and renewable energy sources. However, there are still various mid-size data center installations that can improve their power usage by implementing simple improvements.

SUSTAINED VERSUS SURGE POWER

To maintain sufficiently cool temperatures, data center designers should plan their cooling systems according to the power used and heat generated—this is determined by the volume of network activity and processing the equipment performs, which can vary significantly throughout a normal day. The sustained power is the amount of power required to perform minimal, constant activities. During points of high network activity, more power is required, and in turn, more heat is generated.

When these peak or surge power usages occur for just a few minutes in an hour, the heat load can usually be managed effectively. However, the longer this activity is maintained, more power is required and more heat accumulates. Additionally, many data centers balance their power and cooling requirements by spreading out various network activities and processing runs to decrease surge times. Instead of designing multiple functions to perform simultaneously, some maintenance functions are scheduled for periods of lower power usage.

Improving Data Center Efficiency and Related Impacts

While excessive heat can be detrimental to electrical equipment, power and cooling systems based on peak load requirements will provide more capabilities than needed—and a much higher price tag (capital) than is necessary. Instead, since periods of surge power exceed the sustained power usage, data centers should be designed for an average volume of power usage.

NAMEPLATE VERSUS ACTUAL POWER USAGE

Similar to planning power and cooling equipment in a data center simply around surge power usage results in excessive capacities and cost, so does planning a system according to nameplate requirements. Equipment contains a nameplate with a label listing the amount of power (watts) that it consumes under 100 percent utilization.

However, this quantity is the maximum possible power supply that could be used—not the typical running power. In fact, this amount is rarely, if ever, achieved, as most equipment designers provide more power supply capacity than the equipment can possibly consume as a precautionary measure. In reality, network equipment running at a high percentage of its nameplate value is likely being overburdened, resulting in slow response time and is in need of replacement or reconfiguration.

Since the nameplate value is greater than the actual equipment power usage, adding the values of all equipment contained within a rack to determine that rack's power and cooling requirements results in an inaccurate and extremely high number. As an alternative, some equipment manufacturers provide guidance in typical power and cooling needs, and the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) also provides a listing of typical running requirements. Actual power usage can also be determined with a PDU that has an Amp meter, which provides instantaneous power consumption when multiplied by the incoming voltage. Recorded over time, this provides accurate overall power consumption and is invaluable to understanding total power usage in a data center. Finally, a common industry practice is to de-rate the nameplate by 40 percent to garner an approximate actual equipment operating power.

Positioning the cooling device closer to the heat source will improve your cooling efficiency

Power is directly related to the required cooling, as nearly all power consumed by processors is converted into heat.The better you plan your required operating power, taking into account future expansion needs, the better you can design your cooling system. In addition, the closer the cooling is to the heat source the quicker it can react to the fluctuating heat loads and hence, you avoid the possibility of under cooling your data center.

RAISED VERSUS NON-RAISED FLOOR DATA CENTERS

In raised floor data centers, perforatedfloor tiles are used to distribute cold air - allowing cold air to enter a cold aisle and be used for equipment intake.

Typically, each floor tile is 600 mm x 600 mm (24" x 24") and allows 300 to 850 m³/h (200 to 500 CFM), depending on the percentage of open space and underfloor air pressure.

Open space in percentage is the amount of cutout area that allows airflow through the tile. Older floor tiles usually provide 25 to 35 percent open space, while some newer models feature open spaces approaching 60 percent of the total tile area.

The exact amount of cooling (watts) per tile is determined by the actual airflow and temperature. Typical floor tiles providing 1,500 to 4,000 watts of cooling.

High-performance tiles composed of grated steel can achieve an open space of 56 percent or greater. Although tiles feature more open space, a whole system installed with these may be difficult to support—computer room air conditioner/handler units often do not have enough air pressure beneath the floor to distribute the required CFM throughout the data center.

Some vendors have developed liquid cooled floor tiles or tiles with gills. Such solutions can help to support cold air or higher airflow at a dedicated area in the data center and should therefore mainly be considered when optimizing an existing installation.

Air recycling occurs through the use of Computer Room Air Conditioners, CRAC, or computer room air handlers, CRAH, units. These air conditioners remove the heat from the hot exhaust utilizing either refrigeration or chilled water. The heat is directed outside of the building, and the reconditioned cold air is pumped into the raised floor plenum then up through the perforated tiles and into the cold aisle. Since the raised floor is the plenum for distributing cold air to the data center, the placement of perforated floor tiles is critical.

Improving Data Center Efficiency and Related Impacts

These tiles are the point of entry for all cold air in the data center, and where they are located will determine the cooling capability and overall efficiency of the facility. They must be strategically located to ensure proper cold air distribution.

The exact placement of floor tiles is often set once the rack and equipment are in place, through a process of measurement as well as trial and error. A more sophisticated, alternative method uses Computational Fluid Dynamics (CFD) models to determine optimal tile placement. In this approach, a model of the data center is created with as accurate information as possible. A floor plan is developed, and through modeling simulation, the placement of the floor tiles that will allow ideal airflow and cooling is attained.

Complications can result from perforated floor tiles being positioned either too far from or too close to the CRAC/CRAH unit(s). Floor tiles placed too far from the CRAC/CRAH unit(s) produce little to no airflow.

Conversely, due to the high velocity of air exiting the bottom of a unit, perforated floor tiles located in close proximity to the points at which the bottom of a rack meets the raised floor plenum may be exposed to the Venturi Effect.

The Venturi Effect occurs when a rush of air flowing below the perforated tile causes suction into the flow. Instead of having an upward flow of cold air, a downward draw of the air from above the floor may occur—the opposite of what is needed.

Over the last years there have been significant work provided by various organizations driving energy efficiency in data centers. Nevertheless, in many installations, the following parameters are representative of typical raised floor data centers:

- Plenum air temperature = 13° to 19° C (55˚ to 66˚ F)
- Air temperature increases to approximately 25,5° C (78˚ F) as it moves upward to higher mounted equipment
- Typical ΔT = approximately 20 K
- Hot aisle temperatures = approximately 38° C (100˚ F)

CHALLENGES ASSOCIATED WITH RAISED FLOOR DATA CENTERS

- Difficulty with seismic ratings. In earthquake prone regions, it is almost impossible to expect (99.999 percent up time) availability of the network using a raised floor.
- Raised floors were originally designed before the need for frequent moves, additions or upgrades typically associated with the short operating life of IT equipment. Because of this, it is very difficult to access under-floor cabling to accommodate changes.
- Older locations might not have sufficient headroom for raised floors.
- The area under a raised floor is considered to be an air plenum that must meet special fire codes and include conduit requirements or special polymers and plenum rated cables.

NON-RAISED FLOOR DATA CENTERS

The primary elements of a raised floor design are intended to provide a cool air distribution system, conduits for power and data cabling, and a copper grid for grounding. However, many of these objectives can be achieved via new technologies. In fact, many large data centers now use non-raised floor designs. Along with the removal of raised floors, drop ceilings are also eliminated. Today, non-raised floor data centers can not only provide sufficient airflow and support structures, but they may also eliminate some of the issues associated with raised floor designs.

ADVANTAGES OF NON-RAISED FLOOR DATA CENTERS

- Cabling is located overhead and is easy to access.
- Cleaning the data center floors is much easier, since there is no place for dust to hide.
- Weight capacity concerns are eliminated with non-raised floor data centers. Racks full of equipment weigh as much as 1500kg (3,300 lbs) or more which can be a major issue with raised floor designs.

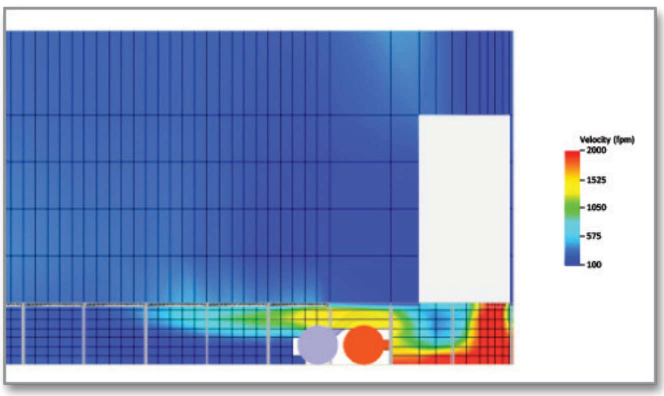

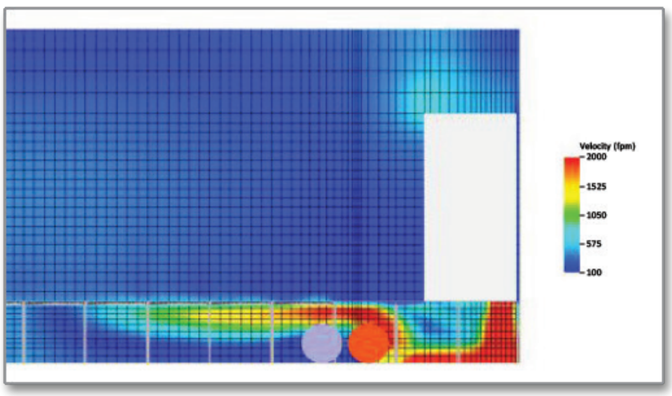

Picture: CFD analysis of same layout with different calculation grids

Improving Data Center Efficiency and Related Impacts

- Costs are lower since there is no raised floor with which to contend. Outside studies suggest raised floors cost approximately €200/m² ($20.00/sq. ft.), plus the added costs associated with power and data cabling.
- Blockages, barriers and obstructions to cold airflows are eliminated. In many instances, perforated tiles deliver substantially less airflow than expected due to blockages and obstructions under the raised floor.

A cooling strategy that delivers cold air to the equipment must be carefully designed

DATA CENTER ENERGY COSTS AND SAVINGS

After the perforated floor tiles' positions are determined, additional steps can be taken to ensure proper airflow in the data center.

Once the cold air is entering over the raised floor you need to ensure proper air separation. The cold air shall be provided were it is required at the intake of network equipment. Exhausted air shall not mix with cold air and cycled back to the cooling unit. By using brushes, gaskets and panels closing potential areas where air short-circuits (mixing of cold and warm air) can be eliminated.

By utilizing a cold aisle containment system, you can separate the cold and hot air areas to improve the efficiency of the raised floor cooling in front of the 19" plane. Also, thermal layers requiring under cooling are minimized resulting in energy savings.

Already introduced in 2006 the PUE is a common metric to measure efficiency in the data centers. The total amount of power used by the data center compared to the power used by the IT equipment.

Main benefit of the metric is to give the operator an indicator on how different activities impact the efficiency.

It can also provide a kind of benchmark for different data centers.

However this need to be seen carefully as geographical location, server utilization and finally details on how power is measured are certainly not always the same.

PUE
(Power Usage Effectiveness) =
Total Facility Power/IT Equipment Power

$$\text{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

A typical data center with a hot aisle / cold aisle rack configuration has a PUE number around 2.4. An efficient data center brings this number down to 2 or even lower. Some purpose-built data centers are claiming PUE numbers of 1.3 or even lower.

DCiE is another metric for measuring data center efficiency. It is similar to PUE but is inverted and shown as a percentage.

A PUE of 2.4 = 41.6%.

DCiE
(Data Center infrastructure Efficiency =
IT Power Used/Total Facility Power)

$$\text{DCiE} = \frac{\text{IT Equipment Power}}{\text{Total Facility Power}}$$

SAVINGS CHART

This chart represents typical savings and is a guideline. Each data center will present many challenges, and the activities shown will vary significantly in their effectiveness and savings. Potential floor space savings and capital construction costs are not included
in calculations.

During the planning stage of a new data center, various layouts can be evaluated based on the required IT equipment and utilization to optimize energy efficiency. Nevertheless, even existing data centers can achieve significant savings by implementing different activities linked to small invests.

HYPOTHETICAL DATA CENTER

- 100 Racks with an IT load of 5 kW per rack (average)
- 3000 sq. ft. (280 m²) raised floor area (5 rows of 20 racks with 4 aisles)
- Total IT load = 500 kW-hr (4,380,000 kW/year)
- Total power used annually 10,512,000 kW
- Energy cost $0.10 / kW-hr
- Total annual energy costs $ 1,051,200.00
- PUE = 2.4 (Typical average data center) DCiE = 41.6%
- Annual cost per rack = $ 10,512
- Annual cost per 1kW of IT load = $ 2,102

Improving Data Center Efficiency and Related Impacts

ACTIVITY PUE ANNUAL POWER COST ANNUAL SAVINGS

Activity	PUE	Annual Power Cost in $	Cumulated annual Saving in $
Data Center typical (Hot aisle/Cold Aisle)	2.40	1.051.200	-
Blanking Panels (In all open RU locations)	2.38	1.042.440	8.760
Floor Brushes (Tile Cutouts)	2.35	1.029.300	21.900
Perforated Tile Placement	2.30	1.007.400	43.800
CRAC/CRAH Unit -Duct Work	2.27	994.260	56.940
Drop Ceiling Return Duct Work	2.23	976.740	74.460
Rack Layout Optimized	2.20	963.600	87.600
Rack Floor Optimized	2.15	941.700	109.500
Containment - Cold Aisle	2.10	919.800	131.400
Containment with in row Liquid Cooling	1.85	831.200	219.000
Liquid Cooling* (Closed Rack System)	1.75	766.500	284.700

*Note: Liquid cooling is a stand-alone solution not cumulative to the other activities.

ABOUT THE AUTHOR

Markus Gerber graduated in business administration in the market and communications research department of the Pforzheim University of Applied Sciences. He has worked for nVent since February 2007. He held the position of project leader in the introduction of a comprehensive service concept, and since January 2009 has adopted a variety of roles in product management for cabinets. Additionally, he can boast a successful track record of working on numerous datacoms projects in the EMEA region.

ABOUT NVENT

At nVent, we believe that safer systems ensure a more secure world. We connect and protect our customers with inventive electrical solutions. nVent is a $2.1 billion global company that provides enclosures, electric heat tracing solutions, complete heat management systems, and electrical and fastening solutions. nVent employs 9,000 people worldwide.

ABOUT ENCLOSURES

Electrical systems come in all shapes and sizes, from massive industrial controls to single components. nVent offers a comprehensive range of enclosures that house these vital assets. Marketed under the nVent HOFFMAN and SCHROFF brands, our enclosures offer two-pronged protection: safeguarding electrical equipment from the operating environment and people from electrical hazards. The nVent SCHROFF brand includes server cabinets, data center cooling solutions, power supplies and subracks and cases.

Contact:

nVent SCHROFF
+49.7082.794.0