

Divide & Conquer

Power Savings from Server to Datacenter

A study by The Department of Engineering at the University of Taiwan to reduce energy costs and improve efficiency in datacenters

By John Parry
Electronics Industry Manager,
Mentor Graphics

Increasingly organizations are choosing to host their data requirements in large purpose-built, energy-hungry datacenters. Datacenters house many racks and a large number of servers where a significant amount of heat is generated from the IT equipment. To remove the heat from equipment so that the electrical components can operate normally, cooling systems must be introduced to provide adequate cooling. Fans, Computer Room Air Handlers (CRAH), and chillers, consume 35%-45% of the total power budget. A big concern for datacenter operators is reducing energy bills and one approach is to improve cooling efficiency.

Figure 1 shows one datacenter room with a single CRAH and 20 racks; each rack is 42U high, and so can house 42 1U-servers. The datacenter includes a raised floor to provide cold air from the CRAH for the IT equipment and a dropped ceiling to draw hot air from the IT equipment into the CRAH. This is a typical datacenter arrangement in which some of the hot and cold air mixes, increasing the room and server inlet temperature. To provide a suitable ambient environment for IT equipment, lowering the temperature may require a higher cooling

capability from the CRAH, which will waste more power.

Divided Zone Partitioning

A recent study by The Department of Engineering at the National Taiwan University explored the idea of a divided zone approach to cooling efficiency. A divided zone partition works by concentrating airflow for key components to avoid airflow bypass and controlling different individual zones independently with the aid of a Fan Speed Control (FSC) for the system.

Mixed airflow challenges can be overcome by cooling airflow path management to improve cooling efficiency and power saving [1]. Some datacenters implement hot and cold aisle containment [2, 3], while Zhou et al. [4, 5] propose adaptive vent tiles that can vary their opening for air flow adjustment. Most hot and cold aisle containment systems encompass and seal off the racks in the same row of a datacenter, but providing adequate cooling performance when the loading of one or more racks is much lower than the others remains an issue.

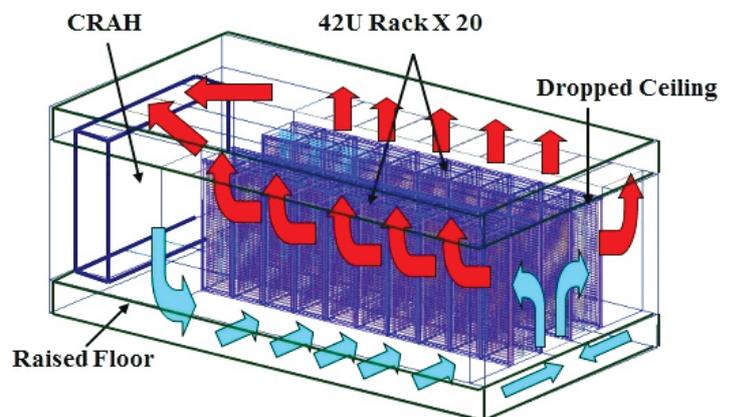


Figure 1. Datacenter configuration and airflow direction

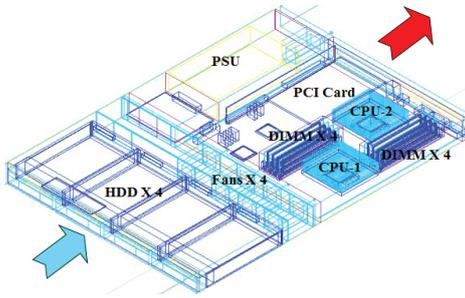


Figure 2. Configuration of the 1U server CFD model and airflow direction

partition shows a significant power saving for IT equipment from the server level up to the datacenter level.

As Simple as One, Two, Three with FloTHERM®

The CFD analysis for the building simulation model is performed with FloTHERM 3D CFD software using a structured Cartesian grid that can be localized and nested to minimize solve times and enable multi-scale modeling for accurate results.

Server-Level Power Savings from a Divided Zone Partition

Figure 2 shows a standard 1U height server CFD model with four fans. Three fans are directed at the CPUs and DIMMs, and one fan towards the PCI card. The PSU includes its own fan, which is located at the rear. We find the thermal solution first to tune the suitable fan curve performance to pass the component thermal specifications at 35°C system ambient. While all the equipment temperatures meet the system thermal requirements, it cannot be assumed that this is the optimum design. Figure 4 shows the airflow distribution. Some airflow does not follow the desired path, providing an opportunity for power savings.

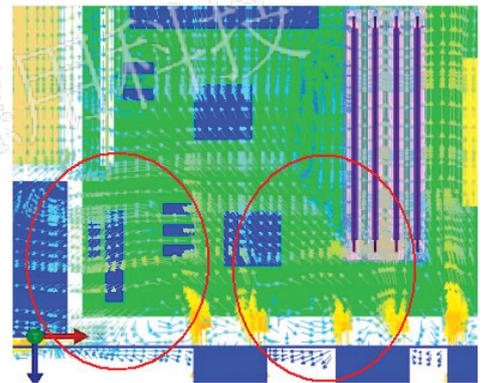


Figure 4. Server system airflow bypass illustration

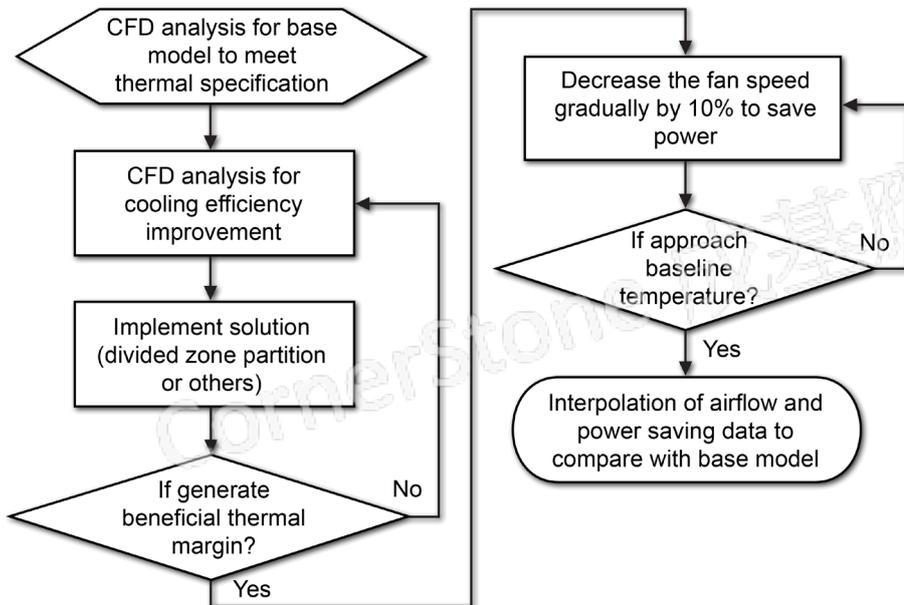


Figure 3. Flow chart of a divided zone analysis for power savings

Servers of the type shown in Figure 2 consisting of Hard Disk Drives (HDDs), Central Processing Units (CPUs), and Dual In-Line Memory Modules (DIMMs), etc., are the main type of IT equipment and can be the target for major power savings in a datacenter, where research [6, 7, 8] into server liquid cooling have shown improvements in cooling efficiency.

In this study, the divided zone method is developed to improve the cooling efficiency for both a server and a datacenter. The effect of a divided zone on airflow management and fan power savings under normal conditions and during a component load change were investigated in detail. Additionally, the utilization of a divided zone

To ensure the design is workable, a fully loaded system is considered at 35°C ambient according to the ASHRAE maximum allowable temperature of the A2 class [9]. First, the thermal solution for the fan and heatsink to meet the fully loaded system requirements that can satisfy all component and device thermal specifications is found for the base model. Second, design optimization is performed to determine whether the divided zone method would improve the cooling. Finally, the solution is analysed to calculate the resulting power savings. The flow chart is shown in Figure 3.

The divided zone partitions were implemented to determine their optimum position for airflow management (Figure 5 overleaf). The resulting effect was that components that generate higher temperatures, such as the CPUs, receive better cooling. The resulting CPU temperature margin allows fan speeds to be reduced, achieving considerable power savings.

The divided zone partition not only saves server fan power directly, but also decreases the system airflow rate requirement, as shown in Table 1. The airflow rate savings can reduce the CRAH blower load. For this case, the divided zone partition can help decrease the server system airflow rate from 59.5 CFM to 51.3 CFM, a 13.8% reduction in the system airflow rate that can further improve datacenter cooling efficiency.

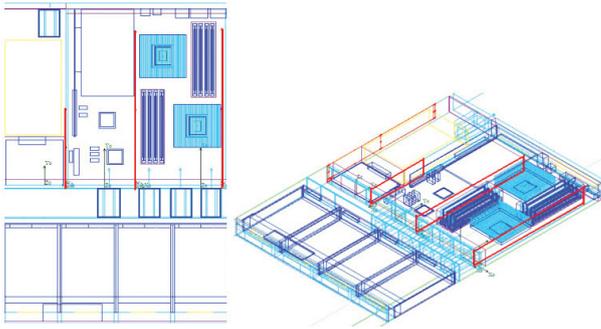


Figure 5. Divided zone partition in the server system

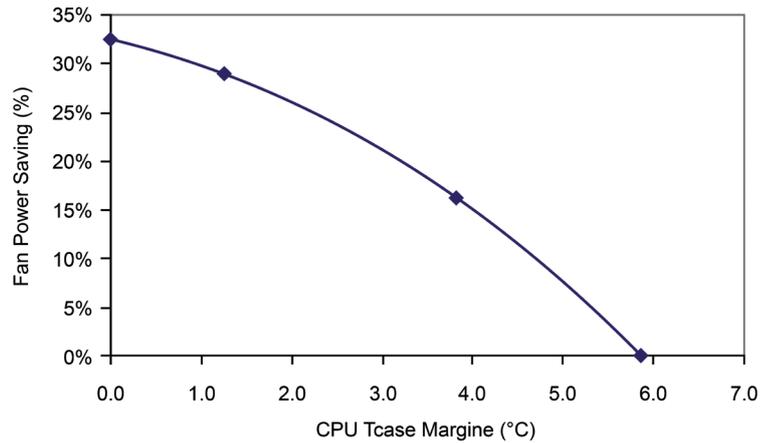


Figure 6. Fan power savings rate vs. the CPU temperature margin of the divided zone partition system

Case	Divided Zone Partition	Fan Speed	CPU2 Tcase (°C)	Fan Power (W)	Fan Power Saving (%)	System Airflow Rate (CFM)	Airflow Rate Saving (%)
1	No	100%	79.4	66.0	-	59.5	-
2	Yes	100%	73.5	66.0	-	59.9	-
3	Yes	90%	75.5	55.2	16.4%	56.1	5.7%
4	Yes	80%	78.1	46.8	29.1%	52.5	11.7%
5	Yes	70%	82.0	39.9	39.5%	48.7	18.1%
6	Yes	76.7%	79.4	44.5	32.6%	51.3	13.8%

Table 1. Simulation result data of the divided zone partition in the server compared with the original model

Server-Level Power Savings with Regional Load Change Condition

The components in a server system are not always at full load, and the load will not be constant, so there will typically be a thermal sensor in the system to detect component temperature variations caused by ambient temperature or load changes. The fan speed can be modulated according to the thermal sensor data by a controller chip in the server. Figure 7 shows the thermal sensor locations. In addition to the ambient sensor placed in the front of the server, there are component sensors placed in CPUs, DIMMs, etc.

The CPUs and DIMMs are controlled by the three fans shown in the right region, the PCI card and chips are controlled by the one fan shown in the middle region, and the PSU is controlled by its own fan in the left region. In this case, the load in the right region was changed to decrease the CPU power from 95 W to 76 W, a reduction of 20%. The CPU was decreased to 71.9°C, providing a temperature margin corresponding to a fan power saving of 31.4%, as shown in Table 2. When the divided zone partition is added, the power savings can be improved to 46.8%. The divided zone partition is not only useful for typical conditions but also for FSC where the local fan is already independently controlled by temperature in a specific area.

Divided Zone Partition Implemented in a Datacenter

The configuration in Figure 1 follows a typical datacenter layout. Earlier studies have shown that air containment separating hot and cold air can be very effective [2, 3]. In this study a simulation model was created for 20 racks with a total of 252kW of power consumption. Small gaps on each side of the rack were included to emulate a true IT equipment environment. Temperatures were set to an ambient of 27°C, the ASHRAE recommended temperature for this class of datacenter [9].

The hot air from the outlet of the racks and

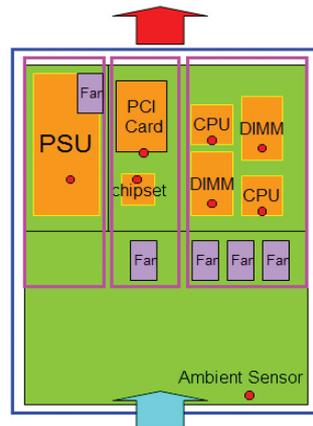


Figure 7. Thermal sensors in the server

the cold air provided to the inlet of the racks are not separated completely, which allows mixing between the hot and cold areas, as shown in Figure 8. To avoid this, some advanced datacenters are implementing air containment, as shown in Figure 9.

At the server level, the component loading is not consistent over a long period of time. At the datacenter level, the rack also experiences different loadings. Although a datacenter with air containment can improve cooling performance, different loadings of the racks in the same row still cause problems. Assume that one of the racks has a zero loading state and that there is no driving fan in the rack. With hot air containment, Figure 10 shows a backflow of hot air from the rear of other racks through the zero-loaded rack to the cold air area. The temperature of the zero-power rack, shown in Figure 11, is seen to be higher than that of other racks.

The partition material could be flexible and transparent plastic sheet for easy access and low cost, with the benefit that every rack can operate independently from the others. From the FloTHERM results, the divided zone has a significant influence on the improvement of the cooling performance of the datacenter under different loading conditions. The study showed that a rack filled with servers with a lower load and correspondingly lower server fan speed resulted in inadequate rack airflow when contained alongside racks containing high-load, high-airflow servers.

The divided zone partition can improve the situation by increasing the airflow rate in the separated region. This can provide better cooling performance for the specific rack and prevent the rack from receiving inadequate airflow, which would lead to an increase in server fan speed and power consumption as the FSC function worked to meet the server's thermal specifications.



Conclusions

A divided zone method has been successfully developed to improve the cooling efficiency for a datacenter. The performance was simulated and investigated under different operating conditions. The major findings are that the divided zone partition can avoid airflow bypass to gain power savings. The partition can save 32.6% of the total fan power consumption and reduce the server airflow rate by 13.8%, reducing the CRAH blower load. For a specific load change case in the server, the FSC function can save 31.4% of the fan power consumption when the CPU load decreases from 95 W to 76 W. Power savings can be enhanced from 31.4% to 46.8% by implementing the divided zone partition with the FSC function. For an advanced datacenter design, the air containment system can avoid the mixing of hot air and cold air to improve cooling efficiency. However, different rack operating load in the same containment region remains an issue. For a 30% load rack case, the implementation of a divided zone

partition in the air containment system can improve the airflow rate by 39% for the fan operating in a server case. The divided zone partition shows a significant power savings for IT equipment from the server level to the datacenter level, making it a good choice for datacenter refits to reclaim lost capacity due to cooling.

References

1. M. Green, S. Karajgikar, P. Voza, N. Gmitter, and D. Dyer, "Achieving Energy Efficient Data Centers Using Cooling Path Management Coupled with ASHRAE Standards2, Proc of 28th IEEE SEMI-THERM Symposium, pp. 288-292, 2012.
2. J. Niemann, K. Brown, and V. Avelar, "Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency", White Paper 135 from Schneider Electric - Data Center Science Center.
3. J. Niemann, K. Brown, and V. Avelar, "Hot-Aisle vs. Cold-Aisle Containment on Data Centers", White Paper 135 from Schneider Electric - Data Center Science Center.

4. R. Zhou, C. Bash, Z. Wang, A. McReynolds, T. Christian, and T. Cader, "Data center cooling efficiency improvement through localized and optimized cooling resources delivery", ASME 2012 International Mechanical Engineering Congress & Exposition, 2012.
5. R. Zhou, Z. Wang, C. E. Bash, C. Hoover, R. Shih, A. McReynolds, R. Sharma, and N. Kumari, "A holistic and optimal approach for data center cooling management", 2011 American Control Conference, 2011.
6. M. Iyengar et al, "Server Liquid Cooling with Chiller-less Data Center Design to Enable Significant Energy Savings", Proc of 28th IEEE SEMI-THERM Symposium, pp. 212- 223, 2012.
7. P. R. Parida, "Experimental Investigation of Water Cooled Server Microprocessors and Memory Devices in an Energy Efficient Chiller-less Data Center", Proc of 28th IEEE SEMI-THERM Symposium, pp. 224-231, 2012.
8. M.M. Ohadi, S.V. Dessiatoun, K. Choo, and M. Pecht, "A Comparison Analysis of Air, Liquid, and Two-Phase Cooling of Data Center", Proc of 28th IEEE SEMI-THERM Symposium, pp. 58-63, 2012.
9. ASHRAE TC 9.9, "2011 Thermal Guidelines for Data Processing Environments - Expanded Data Center Classes and Usage Guidance", ASHRAE, 2011.

Divided Zone Partition	Fan Speed	T CPU2 Tcase (°C)	Fan Power (W)	System Airflow Rate (CFM)	Fan Power Saving (%)
No	100%	71.9	66.0	59.7	-
No	77.8%	79.4	45.3	51.7	31.4%
Yes	60%	79.4	35.1	51.3	46.8%

Table 2. Power savings from the divided zone partition for CPU load change situation

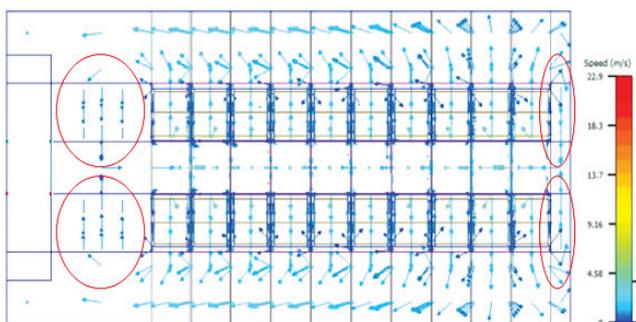


Figure 8. Velocity plot of the simulation results shows the air circulation within the datacenter

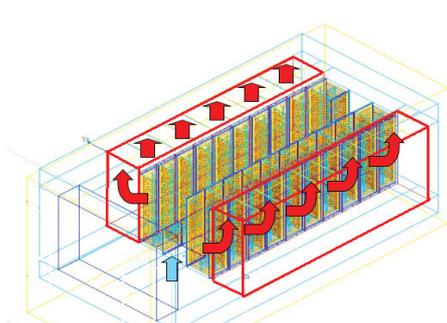


Figure 9. Hot aisle containment within the datacenter

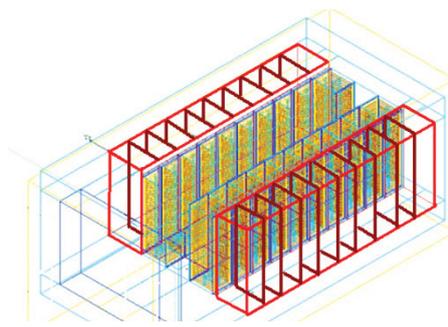


Figure 12. Divided zone partition within the datacenter

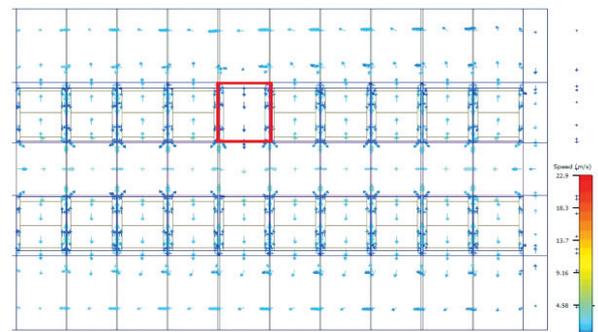


Figure 10. Velocity plot of the simulation result shows that the airflow flows back to the cold air area from the zero loading rack

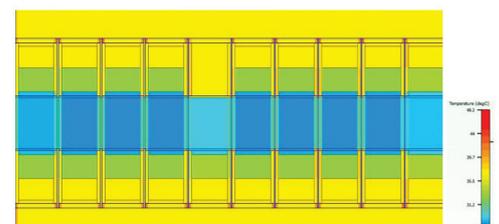


Figure 11. The temperature of the zero power rack is higher than that of other racks



**Mentor
Graphics**

— Mechanical Analysis

This article originally appeared in
Engineering Edge Vol. 2 Iss. 3

Download the latest issue:

[www.mentor.com/mechanical/products/
engineering-edge](http://www.mentor.com/mechanical/products/engineering-edge)

ENGINEERING EDGE

Accelerate Innovation
with CFD & Thermal
Characterization

You might also be interested in...

Toyota: Hybrid Innovation - Thermal Characterization in
Automotive Hybrid Inverter Power Modules

Philips TV's make LEDs Dance: LEDs Transform the Picture
Voxdale & Bernico International: Breaking a World Water
Speed Record

FAdeA Marry Simulation & Test

Eurotecnica Harness the Power of the Sun
How To Guide: Optimizing IGBT Design using T3Ster &
FloTHERM

25 Years of Thermal Analysis

Download the latest issue:

www.mentor.com/mechanical/products/engineering-edge

©2013 Mentor Graphics Corporation,
all rights reserved. This document
contains information that is proprietary
to Mentor Graphics Corporation and may
be duplicated in whole or in part by the
original recipient for internal business
purposes only, provided that this entire
notice appears in all copies. In accepting
this document, the recipient agrees to
make every reasonable effort to prevent
unauthorized use of this information.

All trademarks mentioned in this
publication are the trademarks of their
respective owners.