

How Much Containment Is Enough?

By Ian Seaton
Global Technology Manager
iseaton@chatsworth.com

Published May 2012



CHATSWORTH
PRODUCTS, INC.

800-834-4969

techsupport@chatsworth.com

www.chatsworth.com

How Much Containment Is Enough?

It wasn't that long ago that the central question was, "What is data center hot air containment?" As the term began to gain some traction in the industry, the question became, "What is the value or benefit of data center hot air containment?" Today, we can safely say, despite some lingering attention-seekers, that the concept and value of containment is at least somewhat established in the industry (Figure 1) and the more meaningful discussion may now drift toward determining how much containment is enough. As with the differences between hot aisle containment, cold aisle containment and cabinet-level containment, there is no one-size-fits-all answer to the optimum degree of containment. Architectural environments, business objectives, deployment constraints, and cost-benefit curves will all weigh on intelligent decisions.

Data center containment is no longer a well-kept secret. Data center conferences typically contain one or more presentations on some form of containment and most information technology trade shows will typically have multiple containment solutions displayed. In addition, readily available research by the Lawrence Berkeley National Laboratory and ASHRAE recommend containment. It is specified as a minimum requirement for new data centers in the current draft of California Title 24-2013 Energy Code and it is defined as a desired best practice by the European Data Center Code of Conduct.




Figure 1: Sample CFD representation of data center airflow containment. Note the absence of air mixture or hot spots, allowing cool supply air on the right to remain completely isolated from the warm return air on the left.

Equally compelling arguments can be offered for "a little containment is better than no containment," as well as "maximizing containment maximizes returns."

For the former postulate, a colocation data center in Amsterdam provides convincing evidence. This particular data center underwent a live, online conversion from a water-cooled central chiller plant to six cells of the KyotoCooling air-to-air heat exchange adaptation of energy recovery wheel technology. The conversion included building cold aisle containment structures integrated to a fan wall at one end of each containment aisle delivering air from the KyotoCooling cells. The effectiveness of the cold aisle containment architecture was compromised by the absence of blanking panels in unused rack-mount unit (U) spaces in the server racks; in fact some of the cabinets were only

half full and offered substantial openings between the cold aisles and the hot aisles. The compromised containment resulted in a temperature differential (ΔT) between the hot aisle and the cold aisle of only 10-12°F (5.6-6.7°C), likely less than half the actual ΔT of the air through the IT load.

Nevertheless, because of the remarkable efficiency of the KyotoCooling technology, particularly at a 2N configuration, the resulting air movement energy accounted for less than one-eighth of peak rating. And because the Amsterdam climate necessitated barely 100 hours a year of supplemental mechanical cooling, this “inefficient” design allowed IT load density to increase from 40 watts per square foot to 120 watts per square foot (430 watts per square meter to 1291.6 watts per square meter), improve PUE from 2.5 to 1.15 and save an average of €8000 (\$10,583) a day on cooling energy costs. The economic impact of over-supplying bypass air through the leaky cold aisle containment was insignificant compared to the initial gains and therefore this customer was satisfied with a very low level of containment.¹

Conversely, a case study reported on by Lawrence Berkeley National Laboratory (LBNL) of the National Energy Research Scientific Computing Center in Oakland, Calif. provides an example of a containment implementation that would have benefited from maximizing the level of containment.¹ In this experiment, a cold aisle isolation barrier was built which resulted in a reduction of supply air variation of 26 to 18.5°F (14.4 to 10.3°C). This reduction in mixing between supply and return air resulted in:

1. 75% reduction in fan energy
2. 30-49% increase in sensible cooling tons due to higher return temperatures
3. 3% reduction in pump energy
4. 4% reduction in chiller energy from increase to 50°F (10°C) leaving water temperature (LWT)
5. 21% reduction in chiller energy due to extra water-side economizer hours with 50°F (10°C) LWT

The savings achieved for the cold aisle containment (CAC) pod were extrapolated across the entire data center, resulting in a calculated estimated savings of 740,000 kW hours for fan energy, 110,000 kW hours for pump energy, 210,000 kW hours for chiller plant energy and 1,140,000 kW hours for water-side economization, totaling 2.2 million kW hours, or \$220,000 at \$0.10 per kWh. While that is a noteworthy savings, it is not of a scale such as the initial Dutch data center example, which tempered any motivation to make any additional investments in improvement. As a matter of fact, the LBNL results could also be dramatically improved with an increase in the degree of containment. For example, a total 3°F (1.7°C) supply temperature variation is achievable with a good containment system that includes excellent containment within the server cabinet as well as around it.

By reducing the supply air temperature variation from 18.5°F (10.3°C) to 3°F (1.7°C), additional fan energy can be saved and chilled loop water temperatures can be further elevated to promote greater chiller efficiency and access to more economizer hours. Based on fan affinity laws [$(N_1/N_2)^3 = (P_1/P_2)$], the original LBNL result of 75% reduction in fan energy was produced by a 37% reduction in airflow produced, i.e., $(63/100)^3 = 25/100$. An additional 20% reduction in required airflow production – due to the more complete isolation between supply air and return air and the resultant elimination of the need for some amount of bypass air to maintain a cool curtain in front of the servers – results in a need for only 50.4% of the original airflow requirement and 12.8% of the original fan energy requirement, i.e., $(50.4/100)^3 = (12.8/100)$.

While the LBNL case study reduced the supply air temperature variation to 18.5°F (10.3°C) to allow an increase in the chilled water loop temperature to 50°F (10.0°C), the associated set point was likely producing 58.5°F (14.7°C) supply temperature in order to maintain a maximum 77°F (25°C) server inlet temperature, which would have been the ASHRAE upper recommended environmental threshold at the time of the study. With no more than a 3°F (1.7°C) variation, the supply air temperature can migrate up to as high as 74°F (23.3°C) and still guarantee the same maximum server inlet temperature, and a 64°F (17.8°C) chilled water loop will effectively deliver a 74°F (23.3°C) supply set point. Conservatively assuming a linear slope for the kW per ton chiller operation at different H₂O LWT plot points, the revised chilled water loop temperature would produce a 13.2% savings in chiller plant operation over the baseline.

Finally, based on 2007 hourly bin data, this Oakland, California, data center would have 1200 hours below 45°F Wet Bulb (WB) (7.2°CWB), which would allow a 5°F (2.8°C) approach temperature to the 50°F (10°C) chilled water loop temperature. However, at 59°FWB (15°C), which allows a 5°F (2.8°C) approach temperature for the 64°F (17.8°C) chilled water loop, there would be 7600 hours of water-side economization. The 6.3 times additional hours of economizer cooling translates directly into commensurate chiller energy savings (Figure 2).

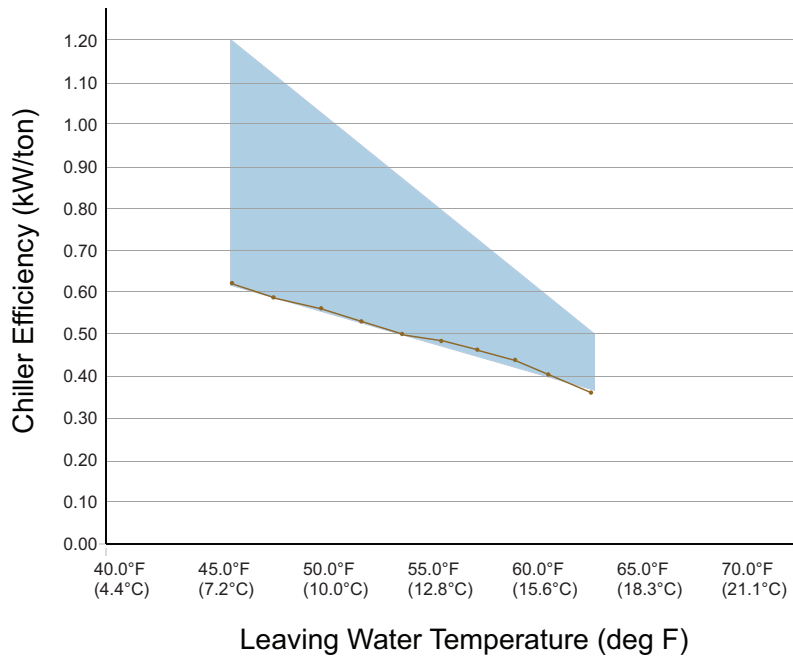


Figure 2: Range of chiller efficiency (kW/ton) as a function of leaving water temperature, with all other parameters held essentially constant (based on ASHRAE Handbook²).

In summary, the reduced variation in server inlet temperatures, resulting from a more effective containment architecture, produces 860 kW hours in total fan energy savings, 412 kW in chiller efficiency savings, and 7.5 million kW hours in economizer savings, for a total savings of 8.9 million kW hours (including original 110,000 kW hours pump energy savings), or 4 times the savings achieved by the LBNL case study containment project. At \$0.10 per kW hour for electrical energy, the improved containment would save \$890,000 from the original baseline or an

additional \$670,000 beyond the savings achieved by the LBNL case study containment improvements. In this instance, more containment obviously is justified.

The above two examples suggest one answer to the question on how much containment is enough. If cooling is practically free at the temperature at which a data center operates with minimal containment, then minimal containment is enough. If, however, higher data center operating temperatures are required to access a greater number of free cooling hours and the non-free cooling hours operating expenses are also temperature-sensitive, then it takes much more containment to reach "enough."

How Much Containment is Available?

A question related to how much containment is enough is the question, "How much containment is there?" In the beginning of data center containment, this was a relatively easy question to answer because you merely had to measure the total volume of "holes" versus the total space and calculate it as a percentage. As containment structures have become more sophisticated and as server loads have introduced more variables, measuring containment has become more complicated.

Containment leaks, which include conductive heat transfer through cabinet skins, tend to be less visible and therefore not amenable to measurement by a ruler or even a caliper. In addition, as IT equipment manufacturers work to improve their energy efficiency, we will see design trends toward higher ΔT 's or high fixed exhaust temperatures, or both, incorporating variable speed server fans to control and reduce that fan energy. For these reasons and for test and validation purposes, it may make better sense to test for containment leakage rates with load banks rather than with IT equipment. With load banks, exact flow rates and ΔT 's can be measured on the work bench and then when the load banks are populated into the containment cabinets, whether it be hot aisle containment (HAC) or CAC or cabinet containment (vertical exhaust duct, chimney, or heat collars), differentials from the pre-determined benchmark can be actually measured.

For example, a six cabinet CAC pod with 20 kW cabinets might have a total airflow rate of 18,600 CFM (31,602 CMH) at a 20°F (11.1°C) ΔT . Any variations in inlet temperatures would represent deviations from the benchmark as would any variations between the supply side of the CAC and the return side of the CAC from the workbench ΔT 's. Let's say that we measured an average 1°F (0.6°C) absolute number deviation (plus = bypass, minus = recirculation and either direction is still leakage, which would either cause an increase in server temperature variation or a decrease in return temperature, with resultant loss in cooling unit efficiency or partial economization where the return temperature marks the useful threshold). The baseline would be $18,600 \text{ CFM} = (3.1 \times W) \div 20$, where $W = 120,000$ watts. The average 1°F (0.6°C) deviation from the baseline would equate to $\text{CFM} = (3.1 \times 120,000)/3$, or 17,714 CFM (30,096 CMH), or a loss of 886 CFM (1505 CMH), or 4.8%. In a particular sample of readings with a 1°F (0.6°C) average and a 0.82 standard deviation, a maximum 3°F (1.7°C) temperature variation would be predicted for somewhere in the system.



Containing Static Pressure

Another key variable in determining the effectiveness of containment is static pressure. The greater the pressure differential (ΔP) between the supply side of the containment and the return side of the containment, the greater the opportunity for increased leakage. When the supply side is positive to the return, that leakage will be bypass and when the supply side is negative to the return, that leakage will be waste air recirculation.

Because Chatsworth Products, Inc.'s (CPI) first market introduction of containment was cabinet containment with the Vertical Exhaust Duct System, we were initially most interested in boundary conditions with the absolute minimum pressure differentials to remove any concerns about the necessity of over-pressurizing a space (Figure 3 and 4). This ensured that the heat could be driven out of the rear of the cabinet and up the Vertical Exhaust Duct or, conversely, eliminated concerns about how the Vertical Exhaust Duct System would function during a loss of air handler fan power. Therefore, we focused on test conditions in the area of 0.003" H₂O (0.75 Pa) column differential between the supply side and the return side. This demonstrated that -0.001" to -0.0015" (-0.25 Pa to -0.37 Pa) was going to be adequate to evacuate 3000 CFM (5097 CMH) of high temperature air out of the rear of the cabinet through the Vertical Exhaust Duct without aid of additional fans. An ancillary benefit of that testing was that it demonstrated the amount of air handler fan energy that could be saved with the cabinet containment system, i.e., production of supply air volume could be cut down to the absolute minimum differential from IT load air demand.




Figure 3: Test example of a Hot Aisle Containment Solution by CPI, with instrumentation locations included for temperature (red) and pressure (green) sensors.

However, with HAC and CAC, over-pressure rather than under-pressure becomes the breaking point of total system effectiveness and efficiency, so the associated testing program needs to drive to higher ΔP 's. To that end, CPI's CAC and HAC standard testing program runs at 0.05" H₂O (12 Pa), and some custom tests have run as high as 0.15" H₂O (37 Pa).




Figure 4: Example of pressure data mapping in a Hot Aisle Containment system.

While tests at higher pressure levels effectively demonstrate the superiority of one containment solution over another, best practices remain that a minimal ΔP between supply and return will always be preferable. Not only will that minimal ΔP translate into fan energy savings, but recent research indicates that servers themselves become a path for significant leakage when ΔP 's are not managed closely.³

The Tate study on leakage in containment environments was inspired by the discovery of lower than expected ΔT 's across cooling coils than would be indicated by known ΔT 's across associated IT loads. What they found was a surprising level of leakage through cabinet mounted servers. For the tested sample of servers running at idle, they found a 23.5% CFM airflow increase at 0.02" H₂O (5 Pa) static pressure, 39.7% at 0.04" H₂O (9 Pa) and 47.8% at 0.05" H₂O (12 Pa) —so a fully configured cabinet could be expected to leak around 400 CFM (680 CMH) of bypass airflow. While the Tate study attributed this bypass airflow problem to both positive pressure inside a CAC and negative pressure inside a HAC, the CAC would seem to be more problematic since the smaller contained space would provide a shorter path to reaching an over-pressurized state. Therefore, the answer to how much containment is enough is not merely found in the efficacy of the containment barrier. It also resides in the total pressure management of the contained environment, particularly with cold aisle containment. Therefore, a complete containment architecture must include an effective pressure differential management system.

Which Containment to Use?

From the very advent of containment, there has been an ongoing debate about which form of containment is the most efficient. The debate, based on equal parts vested vendor interest, intuition and anecdotal evidence, has finally been addressed by a scientific study.⁴ Experiments conducted by Intel and T-Systems in the DataCenter 2020 data center test laboratory in the Munich-based Euroindustriepark found there was no efficiency advantage for one form of containment over another. Therefore a decision regarding containment choice can rest on other architectural and business variables.

The following three paragraphs surface some of the relevant issues with CAC, HAC and Vertical Exhaust Duct Systems.

Cold Aisle Containment

Cold aisle containment (Figure 5) is typically going to be a preferable choice for retrofit applications, especially in situations where significant overhead obstacles exist, such as power and data cable distribution pathways and lighting. According to the Intel T-Systems study, cold aisle containment provided greater thermal ride-through time for high density (17 kW per cabinet) during a cooling system failure. CAC's are typically associated with raised floor, but not necessarily so, as revealed by the fan wall KyotoCooling cell delivery in one of the opening examples and the overhead delivery of the much-publicized NetApp Energy Star data center in North Carolina. CAC is typically deployed a full pod at a time and will require some penetration through the aisle enclosure for fire suppression. It is compatible with row-based cooling and will require effective pressure monitoring and management. A CAC means that the rest of the data center is the hot aisle, which today means a data center ambient temperature over 90°F (32.2°C), and future working spaces approaching 130°F (54.4°C).



Figure 5: Standard Cold Aisle Containment Solution from CPI. In the example above, CPI Aisle Containment Doors are combined with an overhead ceiling and the F-Series TeraFrame® Cabinet System to enclose the cold aisle between adjacent cabinet rows. The overhead ceiling traps cold air in the aisle so that it will be used to cool equipment, instead of the entire room.

Hot Aisle Containment

Hot aisle containment (Figure 6) is usually the easiest integration with fire suppression because even if the pitch doesn't align perfectly with the containment common duct, the lines will be above the duct top and/or ceiling and drops into the containment area do not need to be perfectly vertical.

HAC will be more flexible for accommodating various supply delivery architectures, making a raised floor unnecessary. According to the Intel T-Systems experiments, HAC had better thermal ride-through time for lower density applications. Like CAC, HAC is typically deployed a pod at a time, is compatible with row-based cooling systems and it will create a high temperature work area that is more confined than CAC.



Figure 6: Standard Hot Aisle Containment Solution by CPI. In the example above, CPI's Aisle Containment Doors are combined with the F-Series TeraFrame® Cabinet System and an overhead duct to enclose the hot aisle between adjacent cabinet rows and the overhead duct is used to isolate hot exhaust air from the room.



Figure 7: F-Series TeraFrame® Cabinet System with Vertical Exhaust Duct

Cabinet Containment

Cabinet containment (i.e., Vertical Exhaust Duct, chimney, heat collar) will provide the most comfortable work area, with ambient conditions defined by the supply air set point and the minimal resultant variation, usually a 2-4°F (1.1-2.2°C) range in the mid-to-upper 70s (23-26°C). The deeper cabinets required for cabinet containment (Figure 7) typically increase the row pitch from 14 to 16 feet (4.3 meters to 4.9 meters), with the resultant fire suppression needing to be laid out on pitch of eight feet (2.4 meters), instead of the more typical pitch of 10 feet (3.0 meters). In a larger space, the reduced pitch may add additional fire suppression parallel runs. In addition, thermal ride-through can be either a plus or a minus for cabinet containment. Cabinets with the Vertical Exhaust Duct that merely raise the level of the return air without directly coupling to a suspended ceiling or return air duct will have the best thermal ride-through because of the conductive heat absorption of the extra metal surface area in the data center.⁵ However, with the Vertical Exhaust Duct coupled to the contained return air path, that performance diminishes. If the volumetric space for return air exceeds the volumetric space for supply to help sustain a pressure differential during air movement failure, the ride-through time is

still superior; however, when that ratio reverses, the cabinet containment architecture will be the worst of the alternatives. Like HAC, a raised floor is not required for cabinet containment and it can accommodate multiple supply delivery architectures, e.g., overhead duct, fan wall, etc. Cabinet containment data centers can be deployed one cabinet at a time, rather than one pod at a time, and will have an air movement threshold lower than HAC or CAC due to the effect of pressure inside the rear of the cabinet. For CPI solutions, that threshold is 3050 CFM (5182 CMH), which equates to anywhere between 19.5 kW up to 50 kW per cabinet, depending on server fan design efficiency. Most other chimney cabinets tend to have a much lower threshold and therefore compensate for that low airflow by adding chimney fans.

Finally, equipment that does not breathe front-to-rear or front-to-top often compromises the integrity of the best containment system or, worse yet, establishes an excuse why deploying containment is not going to work and/or be worth the investment and effort. These responses to nonstandard air path IT equipment are ill-informed and wasteful. There is no reason to compromise or even avoid containment due to equipment with sub-optimized airflow paths. For side-to-side breathing equipment, standard equipment cabinets have been available to the market for a couple years now. For front-to-side, side-to-rear, side-to-front, or even rear-to-front breathing equipment, simple rack-mount shrouds and duct assemblies provide a path for integration into fully contained spaces (Figures 8-11).




Figure 8: Adjustable QuadraRack® with Evolution® Cable Management and custom baffle for a large network switch that requires side-to-side airflow.




Figure 9: Inlet duct for top-of-rack switch requiring side airflow to avoid cable blockage of normal airway path.




Figure 10: Apparatus for switch requiring side-to-rear airflow.




Figure 11: Apparatus for switch requiring rear-to-front airflow.

Conclusion

In summary, for most of us, with containment that holds leaks to less than 5% of server airflow demand and keeps server inlet temperature variations at 2-4°F (1.1-2.2°C), we will be able to raise supply air temperatures to a point that maximizes all of the following: efficiency of chiller systems, access to free cooling hours and return air temperatures. These high cooling unit efficiencies and increased partial economization hours allow for a quick recoup on any investment in containment. Nevertheless, for data centers with hyper-efficient free cooling systems in climates where high temperatures seldom, if ever, exceed the approach temperatures to the economization solution, and associated IT densities remain low, containment may not need to be as complete.

For all other applications, the answer to “how much containment is enough” will rely on the definition of two key principles: What are the maximum energy savings you could reach through a containment strategy that has been optimized for airflow, static pressure, leakage, bypass air and temperature variance, and how important it is for you to reach those savings?



Ian Seaton
Global Technology Manager, Chatsworth Products, Inc.

Ian has over 30 years of mechanical and electro-mechanical product and application development experience, including HVAC controller components, automotive comfort environmental controls, and aerospace environmental controls and, for the past 13 years, he has spear-headed Chatsworth Products’ data center thermal management initiatives. He is the working group leader for the rack and cabinet section of BICSI-002-2010, Data Center Design and Implementation Best Practices, and also served on the mechanical working group for that standard. He serves on ASHRAE TC 9.9, has published numerous articles and research papers and has presented at conferences and technical meetings in 12 different countries around the globe. He has bachelor’s and master’s degrees from California Polytechnic State University.



References

- ¹Silicon Valley Leadership Group, "Case Study: Lawrence Berkeley National Laboratory Air Flow Management," 2006.
- ²American Society of Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE), "Best Practices for Datacom Facility Energy Efficiency," second edition, 2009, p.41.
- ³Kennedy, Daniel, "Ramification of Server Airflow Leakage in Data Centers with Aisle Containment," Tate white paper, January 6, 2012.
- ⁴Weidmann, Rainer and Leberecht, Markus, "Data Center 2020: Hot-Aisle and Cold-Aisle Containment Efficiencies Reveal No Significant Differences," DataCenter 2020 whitepaper, 2011.
- ⁵Khankari, Kishor, "Thermal Mass Availability for Cooling Data Centers During Power Shutdown," ASHRAE Transactions, Volume 116, Part 2, 2010.

