

CPI Passive Cooling® Solutions: A Path to Higher Density and Lower Cost

by Ian Seaton



CHATSWORTH
PRODUCTS, INC.

800-834-4969

techsupport@chatsworth.com
www.chatsworth.com

It is no longer news to anyone that the effects of Moore's Law has created computer equipment heat loads that data centers are ill-equipped to effectively manage. In fact, this imbalance between potential heat densities and data center cooling capacities is frequently referred to as a crisis today. Emerging conventional wisdom is beginning to settle on a response to this crisis that either involves dispersing the offending servers to reduce the heat densities or utilizing supplemental liquid cooling to remove the heat.

Chatsworth Products, Inc. (CPI) Passive Cooling® Solutions are a better choice because they reduce energy consumption, have lower construction costs, lower operational costs and meet Tier IV operating requirements. CPI Passive Cooling Solutions focus on controlling airflow by isolating and removing heat from the data center allowing you to maximize the cooling capacity in your space.

Unfortunately, passive air cooling is often excluded as a solution because of several misconceptions:

1. The belief that there is a ceiling to passive air cooling capacities well below today's potential heat load densities.
2. The belief that high-density passive air cooling systems create unmanageable high return air temperatures.
3. The belief that lower acquisition costs for passive air cooling solutions are overshadowed by significantly higher operating costs compared to liquid cooling solutions.

This paper will explore the reasons for these three beliefs and explain how an intelligently engineered passive air cooling solution not only overcomes these mythological obstacles, but leads to a solution that effectively cools heat loads well in excess of 30 kW per cabinet, while reducing energy costs and the resultant carbon footprint of a data center. The basic principle behind the passive cooling solution that will be presented is to use the equipment cabinet not as a box for housing servers, but rather as an architectural feature of the data center that secures the isolation between the chilled supply air and the heated return air. In addition, this paper will explore some basic misconceptions about liquid cooling and establish the superior uptime reliability of a particular passive air cooling approach, over any other means of cooling high-density data center heat loads.

Misconceptions About Passive Air Cooling

Myth: The belief that there is a ceiling to passive air cooling capacities well below today's potential heat load densities.

The basis for the belief that there is a low ceiling to the cooling capacity that passive air cooling can achieve, resides in an easily observable phenomena tied to the relationship between volume of air, heat load and temperature rise.

This relationship is described by the equation:

$$\text{CFM} = 3.1W / \Delta T$$

Where CFM = Cumulative cubic feet per minute of airflow consumed by all equipment in the rack

W = Watts (cumulative heat load of the rack)

ΔT = Temperature rise in degrees Fahrenheit (input air versus output air)



It should be noted that this relationship is frequently misapplied by considering the total temperature rise through the rack. The total temperature rise through the rack is more likely a measurement of the inefficiency of the data center room cooling system than it is of the heat transfer going on inside the server equipment. For example, if the equipment at the bottom of the rack is seeing 60°F input air, each piece of rack-mount equipment is experiencing around a 20°F temperature rise, and the equipment at the top of the rack is seeing 75°F input air. The resultant temperature rise through the rack would be 35°F, but that ΔT would not reflect actual heat transfer because of the higher input temperatures delivered to the heat sources higher in the rack. Based on the relationship of these factors in the equation $CFM = 3.1W / \Delta T$, this confusion could result in either grossly underestimating the air delivery requirement or grossly over-estimating the actual heat load of the rack.

The perceived ceiling to air cooling capability as described by the relationship in this equation is based on how much air can be delivered out of a perforated access floor tile. A standard value for this is typically placed at around 700 CFM, though there are high performance grates available today which can stretch that side of the equation. Regardless, assuming that 700 CFM of chilled air is available through a perforated access floor tile located in front of a rack containing 1 RMU and 2 RMU servers with moderate temperature rises, that floor tile could cool close to 6 kW. That same floor tile could cool nearly 8 kW of a blade server load with higher temperature rises.




Figure 1: Rack heat load with inadequate cooling air supply

Figure 1 provides graphic evidence of the relationship between air supply/consumption, heat load and temperature rise. This example from CPI's thermal test bed shows a 9.1 kW heat load with 580 CFM of chilled air delivered from the access floor plenum through a perforated access floor tile in front of the rack. With servers running at a 24°F temperature rise, the cooling equation ($580 \text{ CFM} = 3.1W / 24$) would predict that we could be delivering approximately 4.5 kW through that access floor tile. The computational fluid dynamics



(CFD) model of the actual experiment confirms that prediction by graphically illustrating how the chilled air is consumed by the bottom half of the cabinet, thereby cooling approximately half of the 9.1 kW heat load. The CFD report also shows what happens to the rest of the equipment in that rack. In an otherwise well managed space with most bypass air neutralized, the only air left for the under-supplied servers is the return air – recycled heated exhaust air. That is the by definition, source of a clinical hot spot and the basis for the myth about the relatively low performance ceiling for any kind of air cooling, especially passive air cooling.

Therefore, any kind of passive air cooling solution for high-density heat loads will need to eliminate the dependency on chilled air from a perforated access floor tile in front of the rack, as well as eliminate the source of hot spots from re-circulated return air. CPI Passive Cooling Solutions accomplishes both of those requirements and will be clearly demonstrated later in this paper.

Myth: The belief that high-density passive air cooling systems create unmanageable high return air temperatures.

The second reason for the belief that passive air cooling is not compatible with high-density heat loads has to do with the idea that high return air temperatures will prevent cooling units from functioning properly. While that may be true, if an IT manager were to tell his facilities manager that he was experiencing high return air temperatures, the facilities manager would likely respond by telling him to keep up the good work, especially in a chilled water environment. The reason for this unexpected response is quite clear. Chilled water computer room air conditioners (CRAC) improve efficiency; that is, they increase cooling capacity with higher return air temperatures.

Cooling Unit	Supply Air Temperature	Return Air Temperature	Cooling Capacity
Stulz CCD-600-CWE	60°F	70°F	7.8 tons
Stulz CCD-600-CWE	60°F	90°F	15.5 tons
Stulz CCD-600-CWE	60°F	105°F	20.7 tons
Stulz CCD-1800-CWE	60°F	70°F	23.0 tons
Stulz CCD-1800-CWE	60°F	90°F	46.0 tons
Stulz CCD-1800-CWE	60°F	105°F	61.3 tons

Table 1: Effect of Return Air Temperatures on CRAC Performance Ratings

Table 1 shows how significantly cooling unit capacity can be increased with higher return air temperatures. However, there are a couple caveats to this proposition. First, there is not a flexible performance curve for direct exchange (DX) cooling units – this capacity bonus is only available with chilled water units. Secondly, there is a limit to how high this return air can be before the performance curve starts on a path of diminishing returns, primarily by raising the supply air temperature.

CPI Passive Cooling Solutions have produced actual measured data center return temperature differences ranging from 30°F up to 55°F. As Table 1 implies, the first course of action is to specify chilled water cooling unit solutions that accommodate wider ΔT between the supply air and the return air, and which, in fact, deliver superior performance at those higher ΔT . At extremely high-densities and at high utilization levels, it is possible to drive return air temperatures beyond the limits of these highly efficient cooling units. In these



instances, the most typical remedy is to deliver a portion of economizer air to the return air to keep it within the limits of the CRAC performance curve. In those rare instances when the return temperature exceeds the cooling properties of economizer air, then a small amount of bypass air will suffice. Since the ideal embodiment of the passive air cooling solution described in this paper includes isolated return air in a ceiling plenum, bypass air may be delivered through an open ceiling grate. An open ceiling grate is a standard part of the solution anyway, allowing a certain amount of bypass air to help regulate pressure in the room. By making the return air plenum large enough (i.e. double the supply plenum) this system will be self-regulating for most ranges of airflow fluctuation. However, it can be fine tuned with air handlers equipped with variable air volume (VAV) fans tied to temperature sensors at a few key server air inlet points.

In summary, this passive air cooling solution does create high return air temperatures, which is good up to a point and then there are simple site management strategies to allow the data center manager to continue reaping the benefits of high ΔT without driving supply air too high.

Myth: The belief that lower acquisition costs for passive air cooling solutions are overshadowed by significantly higher operating costs than liquid cooling solutions.

The third reason for the belief that passive air cooling is not compatible with high-density heat loads is the idea that the lower acquisition cost for air cooled solutions is more than offset by much higher operating costs, compared to closely coupled cooling solutions, particularly liquid cooling solutions that reside either in the rack or adjacent to the rack. The primary basis for this belief is the inherent inefficiency associated with trying to operate standard hot and cold aisle facilities at higher densities and the resultant over-capacities that are typically driven by the need to supply cooling continuously at a worst case level. With the improved efficiency of a complete isolation between supply air and return air and the resultant operating economies associated with that separation, close-coupled systems lose their operating cost advantage and the total cost of ownership clearly favors the well-engineered passive cooling solution.

One often ignored cost figure is facility construction cost. The first issue with considering construction costs is deciding on a meaningful metric. Most metrics will unfairly favor either a high-density deployment or a more dispersed lower density deployment. For example, square foot costs would tend to unfairly favor a high-density project. Cost per kilowatt, however, appears to be a fair way of comparing all types of deployments, and this is the metric that The Uptime Institute uses as part of its description for the four uptime availability tier classifications (see Table 2).



The Uptime Institute Tier Classifications Define Site Infrastructure Performance



Typical Tier Attributes

Typical Tier Attributes	Tier I	Tier II	Tier III	Tier IV
Building Type	Tenant	Tenant	Stand-Alone	Stand-Alone
Staffing	None	1 Shift	1+Shift	“24 by Forever”
Useable for Critical Load	100% N	100% N	90% N	90% N
Initial Build-Out Gross Watts per Square Foot (W/ft ²) (typical)	20-30	40-50	40-60	50-80
Ultimate Gross W/ft ² (typical)	20-30	40-50	100-150 ^{1,2,3}	150+ ^{1,2}
Class A Uninterruptible Cooling	No	No	Maybe	Yes
Support Space to Raised Floor Ratio	20%	30%	80-90+% ²	100+%
Raised Floor Height (typical)	12”	18”	30-36” ²	30-36” ²
Floor Loading lbs/ft ² (typical)	85	100	150	150+
Utility Voltage (typical)	208, 408	208, 408	12-15 kV ²	12-15 kV ²
Single Points-Of-Failure	Many + Human Error	Many + Human Error	Some + Human Error	None + Fire and EPO
Annual Site Cause IT Downtime (actual field data)	28.8 hours	22.0 hours	1.6 hours	0.8 hours
Representative Site Availability	99.67%	99.75%	99.98%	99.99%
Typical Months to Implement	3	3-6	15-20	15-20
Year First Deployed	1965	1970	1985	1995
Construction Cost (+30%) ^{1,2,3,4,5}	\$220/ft ²	\$220/ft ²	\$220/ft ²	\$220/ft ²
Raised Floor Useable UPS Output	\$10,000/kW	\$11,000/kW	\$20,000/kW	\$20,000/kW

¹ 100 W/ft² maximum for air-cooling over large areas, water or alternate cooling methods greater than 100 W/ft² (added cost excluded).

² Greater W/ft² densities require greater support space (100% at 100 W/ft² and up to 2 or more times at greater densities), higher raised floor and if require over large areas, medium voltage service entrance.

³ Excludes land; unique architectural requirements, permits and other fees; interest and abnormal civil costs. These can be several million dollars. Assumes minimum of 15,000 ft² or raised access floor, architecturally plain, one-story building, with power backbone sized to achieve ultimate capacity with installation of additional components or systems. Make adjustments for NYC, Chicago and other high cost areas.

⁴ Costs are based on 2005 data. Future year costs should be adjusted using ENR indexes.

⁵ See Institute White Paper entitled “Dollars per kW plus Dollars per Square Foot Is a Better Data Cost Model than Dollars per Square Foot Alone” for additional information on this cost model.

Table 2: Typical Tier Attributes
(from “Tier Classifications Define Site Infrastructure Performance,” Turner, Seader, Brill, copyright Uptime Institute)



The lowest tier construction costs cited by Turner et.al., run at \$10,000 per kilowatt, compared to a project using the passive cooling solution described in this paper at \$6,000 per kilowatt, deployed at 500 watts per square foot gross. Since the construction costs in Table 2 are for lower-density “legacy” air-cooled data centers, one would assume some proportionate increase in construction costs for close-coupled higher-density solutions where additional plumbing would need to be run into the data center space.

Operating costs will include many factors, such as illumination, which will not vary based on the cooling strategy, but energy costs for chillers, heat exchangers, condensers, fans and UPS back-up for any of these sub-systems in higher tier applications are all subject to variation based on the cooling strategy deployed. In addition, there are less obvious operating costs associated with close-coupled systems regarding their fail-over options. For example, some “near”-coupled systems rely on proximate continuously running surplus cooling to assure uptime during service or failure. Not only is this surplus an operating expense, it is also an acquisition and related construction cost not associated with the high-density passive cooling solution. Other solutions are so closely coupled that there is no practical path for redundancy, so other more creative fail-over strategies must be devised. For example, with liquid cooled cabinets, the fail-over is frequently a sensor system (which, by the way, is also a potential point of failure) that recognizes the system failure and releases the door to let in air cooling from the room. Obviously, the system assumes the total required cooling load at the cabinet, plus some level of additional cooling delivered at the room level.

Fans generating CFM are another energy consumer in the cooling equation. Some close-coupled cooling systems, because of the tight quarters inside a cabinet and the effect of surface resistance on airflow, have one set of fans blowing cold air into the cabinet and another set of fans pulling the return air into the heat exchanger. Such topologies, therefore, would require at least 3X CFM, plus whatever additional was required for redundancy surplus. Room air handlers, on the other hand, employ the same fan impeller for delivering supply air and drawing in return air, so the base for the room will be 2X CFM, plus any surplus created by reductions of demand from the planned load that drove the capacity plan. These inefficiencies can be expected, but they should never come close to 3X CFM. Furthermore, central air handler fans with variable air volume controls tied to equipment in-take temperatures provide an efficient redundancy back-up system. When these fans are run in parallel, a higher number of low volume fans will consume appreciably less energy than a lower number of fans pushing a higher amount of air. For example, three 30,000 CFM air handlers working at 2/3 capacity and producing a total 60,000 CFM airflow will actually consume about half the electricity as two air handlers moving the same 60,000 CFM. These economies make designing in adequate surplus for handling fail-over situations a logical, though somewhat counter-intuitive choice.

Close-coupled systems, without any redundancies, would theoretically have a one-to-one relationship between heat load and cooling supply; however, the fail-over requirements arising from that close-coupling neutralize any such advantages. In addition, the passive cooling solution with complete isolation between supply air and return air eliminates temperature variations in the room and allows the data center operator to bring a higher temperature supply air into the room. The resultant supply air can be up to 20°F warmer than is typically delivered into data centers that must combat mixing with higher return air temperatures. That extra 20°F opens up a significantly larger number of hours in most geographic areas for reaping extra free cooling benefits from economizers. Table 3 shows the economizer benefits from complete separation between data center chilled supply air and heated return air for both spaces with CRACs in the data center and with external water chilled central air handlers. The available free hours will obviously vary from area to area, but it is clear that either cooling approach with a hot aisle containment strategy will deliver many more economizer hours



than would be available at temperatures below 37°F. The “free cooling” hours are actually not completely free – typically the air movement fans will consume somewhere around 10% of the total HVAC system electrical load. Therefore, in communities like Denver where the wet bulb temperature may not exceed the economizer operating temperature, this solution could reduce total data center annual cooling costs by 90%.

	TIA-942 Best Practices (a)	Contained Hot Aisles (b)	Water-Cooled Central Air Handler (c)
Delivered Air	68-77°F	68-77°F	68-77°F
Supply Air	52-55°F	72-75°F	72-75°F
Water	42°F	62°F	62°F
Economizer	37°F	57°F	75°F
Free Hours	?	?	?

Table 3: Economic Benefits of Hot Air Containment

a.) Well deployed hot aisle and cold aisle separation with water cooled CRACs and water-side economizer

b.) Physically isolated hot aisle separation with water cooled CRACs and water-side economizer

c.) Physically isolated hot aisles with water cooled central air handler with VAV fans, variable speed pumps and evaporative air economizer

While the economic benefits of hot air containment are most dramatic for the actual energy savings for running the HVAC equipment, the benefits are not limited to the HVAC electric bill. A study conducted by the McKinstry Company of eight different geographic regions around the U.S. reveals a more complete picture of the economic benefits dependent on a complete separation between supply air and return air. Table 4 shows the results for a sample location, but the average for all the areas studied reveals additional significant economic benefits:

- 74% reduction in critical refrigeration tonnage
- 35% reduction in peak water use
- 44% reduction in tank that would be required for 24 hours water storage
- 74% reduction in thermal tank volume to allow generators to restart
- 65% reduction in HVAC load on generator, utility distribution
- 78% reduction in three phase component connection requirements
- 89% reduction in floor space requirements for indoor HVAC equipment
- 63% reduction in cost for HVAC water
- 49% reduction in maintenance costs

The hot air containment model also essentially eliminates the regular activity of monitoring hot spots and making adjustments to the room to balance cool air distribution. Furthermore, the containment model, in conjunction with central chilled water air handlers and supply air feedback sensors driving variable air volume fans, as indicated in the HVAC Options Annual Cost Comparison graph in Table 4, delivers significant energy savings at start-up low utilization levels as well.



**Dallas water-cooled CHW plant
Serving 5 Megawatt Tenant
Mechanical System Options**

10/27/06

Using Fort Worth Weather Data
Using Local electrical rates;
basically 9.5¢/kWhr
with (average) \$6.90/kWh demand
charge,
water \$4.72 per CCF

Jeff Sloan, McKinstry Co.
jeffs@mckinstry.com

Key Variables

Owner's criteria:	N+1 redundancy	Percent Populated	100%
Rack heat density	180 W/ft²	Demand kW of racks:	4,500 kW
Area of racks and aisles:	25,000 ft²	Additional load for UPS, PDUs	405 kW
cold aisle temperature:	77 °F	minimum % rh in cold aisle:	30%
Minimum OSA:	900 cfm	maximum %rh in cold aisle:	80%

Base	Option A	Option B	Option C
Chilled Water CRAC units with Water Economizer	Chilled Water Air Handlers with Air Economizer, Steam Humidifier	Chilled Water Air Handlers with Evaporative Air Economizer	Chilled Water Air Handlers with Evaporative Air Economizer
Open Hot-Cold Aisles	Open Hot-Cold Aisles	Open Hot-Cold Aisles	Separated Hot-Cold Aisles

First cost factors - Mechanical contractor

Critical refrigeration tonnage	1473 tons	1502 tons	1510 tons	916 tons
Airflow quantity, fully populated:	550,000 cfm	530,000 cfm	530,000 cfm	440,836 cfm
Peak (hot day) water use:	49 gpm	50 gpm	50 gpm	32 gpm
Tank that would be required for 24 hours water storage:	70,000 gallon s	70,000 gallon s	70,000 gallon s	50,000 gallons
Thermal tank volume to allow generators, chillers to restart:	180,000 gallon s	180,000 gallon s	180,000 gallon s	110,000 gallons

First cost factors - Electrical Contractor

HVAC load on generator, Utility distribution:	1765 kW	1887 kW	1922 kW	1166 kW
Number of 3ø Components to Connect:	43 items	21 items	17 items	12 items

First cost factors - General Contractor

Indoor HVAC Equipment space:	3599 ft²	1945 ft²	1953 ft²	916 ft²
Need for access floor, suspended ceiling:	Both recommende d	Both recommende d	Both recommende d	Neither recommended
HVAC weight on roof:	None	CHW airhandlers, duc t	CHW airhandlers, duc t	CHW airhandlers only

Total Annual Operating Cost: (100% Populated)	\$5,529,622	\$5,504,726	\$5,389,506	\$4,776,153
Computers, UPS, and Transformer Electrical Cost:.	\$4,487,143	\$4,487,143	\$4,487,143	\$4,487,143
HVAC Electrical Cost	\$968,622	\$955,233	\$849,352	\$256,489
HVAC Water Cost	\$39,649	\$35,619	\$28,098	\$12,092
Maintenance Cost	\$34,208	\$26,731	\$24,913	\$20,429

Reliability factors

Method for avoiding hot spots:	frequent adjustment o f floor & RA dampers	frequent adjustment o f floor & RA dampers	frequent adjustment of floor & RA dampers	all adjusted automatically
Method to provide thermal storage:	outdoor insulated chilled water tank	outdoor insulated chilled water tank	outdoor insulated chilled water tank	outdoor insulated chilled water tank

HVAC Options Annual Cost Comparison




Table 4: Effects on Total Cost of Ownership of Complete Hot Aisle Containment From a study of eight U.S. geographic regions by McKinstry Company

While the green field benefits of the total solution associated with complete hot air containment are clear and compelling, the containment model also provides energy saving economic benefits to a deployment in an existing space. Figure 2 shows a data center layout after being modeled with ducted exhaust return air containment. The two short horizontal rows of cabinets hold an average load of 8 kW and the two longer

horizontal rows hold higher loads, averaging around 16 kW. The vertical rows are either low-density cabinets (2 to 2.5 kW each) or connectivity racks. The 8 kW and 16 kW cabinets were all ducted into the suspended ceiling return air plenum. Ceiling grates were added above the hot aisles for the low-density cabinets to keep most of that exhaust air from circulating through the room. The result was to eliminate serious hot spots in the high-density cabinets and to drive up the return air temperature high enough to raise the CRAC efficiency to the point where the design was able to reduce the number of CRAC units from 12 to seven. (Two in the layout in Figure 2 are dormant for back-up).




Figure 2: Data Center Layout with Removed CRAC Units

The resulting total reduction in cooling capacity documented in Table 5 reduced the HVAC energy requirement for this space by 700,000 kW hours per year, prior to accounting for any additional economizer benefits. The annual electricity cost savings for this space with an actual heat load of 800 kW is over \$60,000, plus the additional capital savings for the reduction in CRAC units versus the extra costs for vertical exhaust ducts on four rows of cabinets.

	Free Space Return Air		Contained Return Air Path	
	Quantity	Total Capacity	Quantity	Total Capacity
Stulz CCD-600-CWE	2	40 tons	0	0 tons
Stulz CCD-900-CWE	7	210 tons	3	90 tons
Stulz CCD-1800-CWE	2	60 tons	3	90 tons
Stulz CCD-2300-CWE	1	40 tons	1	40 tons
	Total Capacity	350 tons		220 tons

Table 5: Cooling Capacity Savings from Contained Return Air Path

In summary, we have shown that a low ceiling to the cabinet heat load cooling limit for air cooling can be overcome by eliminating dependence of cooling capacity from specific access floor tiles and removing the effects of mixing return and supply air. Further, we have shown both the benefit of driving higher return air temperatures as well as simple strategies for keeping those temperatures within the limits of the cooling unit's capacity performance curves. Finally, we provided evidence that acquisition costs, plus construction costs, plus operating costs, combine for a lower cost of ownership for high-density passive air cooling solutions over alternative solutions for high-density cooling.

CPI Passive Cooling® Solutions

The remainder of this paper will explain how the elements of CPI Passive Cooling Solutions work together to deliver this performance.

The basic principle for the effectiveness of high-density passive air cooling is for the server cabinet to become the architectural feature of the data center that secures the isolation between the supply air and the return air, so that resultant isolation frees the cabinet from its dependence on specific perforated or grated floor tiles to cool its heat load.

Neutralizing re-circulation of heated exhaust air within the cabinet is a critical part of the total solution. Since the server fans on the in-take side are creating low pressure pockets, air will naturally tend to move toward those spots, including the hot exhaust air from the rear of the cabinet. Most of us today realize the importance of using blank filler panels in all unused rack-mount spaces to control such re-circulation. CPI Passive Cooling Solutions include low cost tool-free Snap-In Filler Panels to facilitate deploying this part of the solution. However, a frequently overlooked source of internal hot air re-circulation is the opening around the perimeter of the equipment mounting area.




Figure 3: Path of Hot Air Re-Circulation Around Perimeter of Equipment Mounting Area

Figure 3 is a CFD model of an actual test bed and shows the path of heated exhaust air along the side of the cabinet between the side panel and the mounted equipment. That re-circulated air is ingested by the server in-take fans in the front of the cabinet, along with the cooled air delivered by the perforated access floor tile in front of the cabinet. In situations where cabinet front doors are not required for either security or aesthetics reasons, some cabinets can secure this path of re-circulation by merely having the equipment mounting rails installed in the full forward location where they would adjoin the cabinet frame and thereby provide an effective isolation barrier between the front and rear of the cabinet. However, where doors are required and features of the servers extend beyond the server face, equipment mounting rails require some increment of



set-back from the cabinet frame and therefore some other means of securing this re-circulation path is required. The CPI Passive Cooling Air Dam Kit secures this path.



Figure 4: Air Dam Kit Prevents Hot Air Re-Circulation Around the Perimeter of the Equipment Mounting Area.

Securing the internal path of return air re-circulation is only the first step. Where cabinet heat loads are underserved by cold aisle air delivery schemes, as they always will be by definition in high-density situations, servers will pull make-up air from the room at large. To the degree that the room has followed best practices by eliminating bypass air, that make-up air will be heated exhaust air. Figure 1 at the beginning of this paper clearly shows the impact of such re-circulation and represents the source of textbook clinical hot spots in the data center. Therefore, the complete removal from the room of the return air provides the solution for maintaining this isolation.

The CPI Passive Cooling high-density solution includes a solid rear door with a door seal gasket, an arced Airflow Director (functionally, a turning vane) and a Vertical Exhaust Duct to complete the isolation of supply air from return air. The gasketed solid rear door prevents heated exhaust air from escaping out the back of the cabinet. However, merely containing that hot air is not sufficient by itself. A reasonably expected result of containing server exhaust air with a solid rear door would be a static pressure build-up in the rear of the cabinet. This could actually create back pressure into the server fans and thereby diminish their air moving performance. To alleviate this potential, the solution includes the Airflow Director that mounts in the bottom rear area of the cabinet.



Figure 5: Airflow Director



The patent-pending Airflow Director functions as a turning vane and catches the exhaust air from the lowest mounted server in the cabinet and turns it upward, thereby preventing any turbulence and resultant static pressure build-up.

The patent pending Vertical Exhaust Duct mounts on the rear half of the top of the cabinet and directs return air out of the cabinet into a drop ceiling plenum, return air duct or high space in the room where mixing with the supply air is minimized. The ideal implementation, however, is to completely remove the exhaust air from the room into a drop ceiling plenum. The Vertical Exhaust Duct bridges from the data center to the return air space and exploits the pressure differentials between the room and the plenum. These pressure differentials assure all flow is from the room, through the servers and into the return plenum.



Figure 6: Vertical Exhaust Duct

The intelligent application of these basic principles of the physics of fluid dynamics creates a high performance passive cooling system that is not reliant on fans or any other active component in the cabinet. Freedom from fans is not merely freedom from potential single points of failure that, by definition, prevent an application from achieving Tier III or Tier IV classification for uptime availability. Fans are actually counter-productive even when they are working. Top mounted fans, by delivering heated exhaust air into the space directly above the cabinet, tend to break down whatever isolation between supply air and return air that can be maintained in standard hot aisles and cold aisles. The result of this breakdown is either re-circulation causing hot spots in the upper area of the cabinet or the need for a lower set point on the HVAC system to keep the entire space colder and thereby negatively impact the total efficiency of the cooling system. In addition, with today's higher densities, the airflow of cabinets will exceed the air movement capacity of most cabinet fans. For example, four chassis of blade servers in a single cabinet will consume and push through anywhere from 1800 CFM to 3000 CFM, thereby exceeding the capacity of cabinet fans to the point that the fans can actually be inhibitors to airflow.

In conclusion, CPI Passive Cooling Solutions have performed adequately in active data centers up to 23 kW per cabinet and have been tested successfully up to 32 kW. This performance level is achieved through the application of sound, but relatively elementary physics principles. Not only does this solution set provide a viable option for cooling higher densities, but it removes any uncertainties associated with potential points of failure, and provides the basis for being able to take advantage of significant energy savings associated with higher HVAC efficiencies and access to greater economizer hours.

