

Ten Cooling Solutions to Support High- Density Server Deployment

By Peter Hannaford

White Paper #42



Executive Summary

High-density servers offer a significant performance per watt benefit. However, depending on the deployment, they can present a significant cooling challenge. Vendors are now designing servers that can demand over 40 kW of cooling per rack. With most data centers designed to cool an average of no more than 2 kW per rack, innovative strategies must be used for proper cooling of high-density equipment. This paper provides ten approaches for increasing cooling efficiency, cooling capacity, and power density in existing data centers.

Introduction

Blade servers use less power than traditional (i.e. 1U) servers. However, when they are housed compactly in a single rack (**Figure 1**), the increased power required and heat dissipated lead to the creation of hot spots in some data centers. It is unlikely that a traditional data center, with computer room air conditioning (CRAC) units supplying airflow through a raised-floor plenum, will be able to deliver a cooling capacity of more than 3 kW to any rack, regardless of quantity and capacity of CRAC units and floor depth. The maximum cooling capacity delivered to a single rack will be even less for rooms with flooded-aisle air distribution over a non-raised floor.¹ Rack-based cooling solutions are now appearing which take cooling capability to levels well in excess of 10 kW per rack.

Figure 1 – Examples of server compaction



In designing the cooling system of a data center, the objective is to create a clear path from the source of the cooled air to the intakes of the servers. Likewise, a clear path needs to be created from the hot exhaust air of the servers back to the return air duct of the CRAC unit. There are, however, a number of factors that can adversely impact this objective.

This paper describes the main factors that reduce the operating efficiency and power density capability of a data center and presents ways to avoid these problems. It also gives some solutions that allow data centers to cope with demands well in excess of the design capacity without major reconstruction.

The following ten sections describe solutions that address the root causes of cooling inefficiency and under-capacity. The ten solutions are presented in rank order, with the simplest and most cost effective presented first. If the objective is to reach data center power density in excess of 6 kW per rack over a sustained area, then going directly to more extreme solutions such as # 9 or # 10, bypassing the early solutions, may be appropriate.

¹ For more information on air distribution architectures see APC White Paper #55, "Air Distribution Architectures for Mission-Critical Facilities."

1. Perform a “health check”

Just as an automobile benefits from regular servicing, a data center needs to be kept operating at peak efficiency to maintain the business processes it supports and to prevent future problems. Before embarking upon expensive upgrades to the data center to deal with cooling problems, certain checks should be carried out to identify potential flaws in the cooling infrastructure. These checks will determine the health of the data center in order to avoid temperature-related IT equipment failure. They can also be used to evaluate the availability of adequate cooling capacity for the future. The current status should be reported and a baseline established to ensure that subsequent corrective actions result in improvements.

A cooling system checkup should include these items:

- **Maximum Cooling Capacity.** If there isn't enough gas in the tank to power the engine then no amount of tweaking will improve the situation. Check the overall cooling capacity to ensure that it is not exceeded by the IT equipment in the data center. Remember that 1 Watt of power consumed needs 1 Watt of cooling. Excess of demand over supply will require major re-engineering work or the use of self-contained high-density cooling solutions described later in solution # 10.
- **CRAC (computer room air conditioning) units.** Measured supply and return temperatures and humidity readings must be consistent with design values. Check set points and reset if necessary. A return air temperature considerably below room ambient temperature would indicate a short circuit in the supply air path, causing cooled air to bypass the IT equipment and return directly to the CRAC unit. Check that all fans are operating properly and that alarms are functioning. Ensure that all filters are clean.
- **Chiller water/ condenser loop.** Check condition of the chillers and/or external condensers, pumping systems, and primary cooling loops. Ensure that all valves are operating correctly. Check that DX systems, if used, are fully charged.
- **Room temperatures.** Check temperature at strategic positions in the aisles of the data center. These measuring positions should generally be centered between equipment rows and spaced approximately every fourth rack position.
- **Rack temperatures.** Measuring points should be at the center of the air intakes at the bottom, middle, and top of each rack. These temperatures should be recorded and compared with the manufacturer's recommended intake temperatures for the IT equipment.
- **Tile air velocity.** If a raised floor is used as a cooling plenum, air velocity should be uniform across all perforated tiles or floor grilles.
- **Condition of subfloors.** Any dirt and dust present below the raised floor will be blown up through vented floor tiles and drawn into the IT equipment. Under-floor obstructions such as network and power cables obstruct airflow and have an adverse effect on the cooling supply to the racks.

- **Airflow within racks.** Gaps within racks (unused rack space without blanking panels, empty blade slots without blanking blades, unsealed cable openings) or excess cabling will affect cooling performance.
- **Aisle & floor tile arrangement.** Effective use of the subfloor as a cooling plenum critically depends upon the arrangement of floor vents and positioning of CRAC units.

For a more detailed description see APC White Paper #40, “Cooling Audit for Identifying Potential Cooling Problems in Data Centers”.

2. Initiate a cooling system maintenance regime

The Uptime Institute² has reported that it found operational deficiencies in more than 50% of data centers visited. Although collectively labeled “poor cooling,” some were caused by inadequate or poorly executed maintenance regimes.

Among deficiencies discovered were:

- Dirty or blocked coils choking airflow
- Undercharged DX systems
- Incorrectly located control points
- Uncalibrated or damaged sensors
- Reversed supply & return piping
- Faulty valves
- Faulty pumps
- Pumps running unnecessarily
- Free cooling systems not initiated

Regular servicing and preventive maintenance is essential to keeping the data center operating at peak performance. If the system has not been serviced for some time then this should be initiated immediately. A regular maintenance regime should be implemented to meet the recommended guidelines of the manufacturers of the cooling components. Contact your maintenance company, HVAC consultant, or APC for more details.

² www.upsite.com

3. Install blanking panels and implement cable management regime

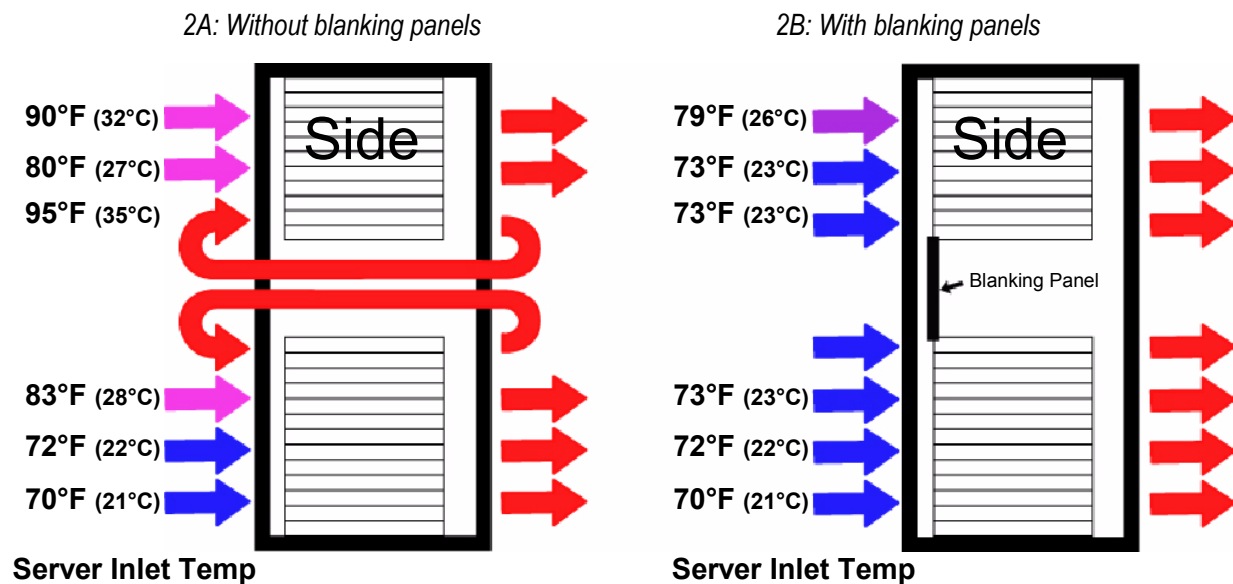
Unused vertical space in rack enclosures causes the hot exhaust from equipment to take a “shortcut” back to the equipment’s intake. This unrestricted recycling of hot air causes the equipment to heat up unnecessarily.

An HP server installation guide³ states:

CAUTION: Always use blanking panels to fill all remaining empty front panel U-spaces in the rack. This arrangement ensures proper airflow. Using a rack without blanking panels can result in improper cooling that can lead to thermal damage.

This advice is often ignored and is a major cause of overheating problems. Installation of blanking panels prevents cooled air from bypassing the server intakes and prevents hot air from recycling. **Figure 2** shows the effect of blanking panels.

Figure 2 – Effect of installation of blanking panel on server air inlet temperature⁴



More information can be found in APC White Paper #44, “Improving Rack Cooling Performance Using Blanking Panels”.

³ HP Proliant DL360 Generation 3 Server Installation Guide, www.hp.com

⁴ APC laboratory experiment – APC White Paper #44, “Improving Rack Cooling Performance Using Blanking Panels”

Airflow within the rack is also affected by unstructured cabling arrangements. When IT equipment is increasingly packed into a single rack, new problems with cable management arise. **Figure 3** illustrates how unstructured cabling can restrict the exhaust air from IT equipment.

Figure 3 – Example of unstructured cabling



Unnecessary or unused cabling should be removed. Data cables should be cut to the right length and patch panels used where appropriate. Power to the equipment should be fed from rack-mounted PDUs with cords cut to the proper length. More information on rack accessories to solve cabling problems can be found on the APC website, www.apc.com.

4. Remove under-floor obstructions and seal floor

In data centers with a raised floor the subfloor is used as a plenum, or duct, to provide a path for the cool air to travel from the CRAC units to the vented floor (perforated tiles or floor grilles) tiles located at the front of the racks. This subfloor is often used to carry other services such as power, cooling pipes, network cabling, and in some cases water and/or fire detection & extinguishing systems.

During the data center design phase, design engineers will specify the floor depth sufficient to deliver air to the vented tiles at the required flow rate. Subsequent addition of racks and servers will result in the installation of more power and network cabling. Often, when servers and racks are moved or replaced, the old cabling is abandoned beneath the floor. This is especially true for co-location and telehousing facilities with high levels of client turnover. Air distribution enhancement devices such as the one shown in **Figure 12** can alleviate the problem of restricted airflow. Overhead cabling can ensure that this problem never even occurs. If cabling is run beneath the floor, sufficient space must be provided to allow the airflow required for proper cooling. Ideally, subfloor cable trays should be run at an “upper level” beneath the floor to keep the lower space free to act as the cooling plenum.

Missing floor tiles should be replaced and tiles reseated to remove any gaps. Cable cutouts in the floor cause the majority of unwanted air leakages and should be sealed around the cables using grommets (**Figure 4**). Tiles with unused cutouts should be replaced with full tiles. Tiles adjacent to empty or missing racks should also be replaced with full tiles.

Figure 4 – Cable cutout grommet



Sealing Cable Cutouts

Cable cutouts in a raised floor environment cause the majority of unwanted air leakages and should be sealed. Based on measurements at multiple data centers, 50-80% of valuable conditioned air is not reaching the air intake of IT equipment due to these unsealed floor openings. This lost air, known as bypass airflow, contributes to IT equipment hotspots, cooling inefficiencies, and increasing infrastructure costs.

Many sites, believing that inadequate cooling capacity is the problem, respond to overheating by installing additional cooling units. One alternative to minimize the cost of additional cooling capacity is to seal cable cutouts. The installation of raised floor grommets increases static pressure under a raised floor. Cool air delivery through perforated tiles and floor grates can also be improved. Sites now can optimize the effectiveness of their existing cooling infrastructure and manage increasing heat loads.

5. Separate high-density racks

When high-density racks are clustered together, most cooling systems become ineffective. Distributing these racks across the entire floor area alleviates this problem. The following example illustrates the effectiveness of this strategy.

Data center design characteristics:

Raised floor area: 5,000 ft² (465 m²)

Raised floor depth: 30 inches (762 mm)

UPS load: 560 kW

Average rack space: 1,250 ft² (116 m²)

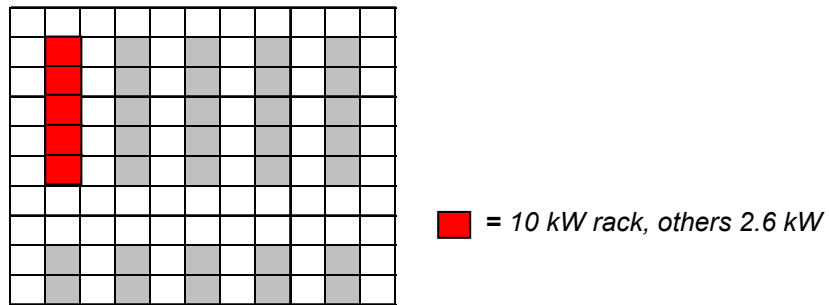
Rack quantity: 200

Average data center power density: 112 watts / ft² (1,204 watts / m²)

Average power density per rack: 2,800 watts

Allowing for aisle spaces and CRAC units, and making the assumption that racks occupy one-quarter of data center floor space, the average rack density would be 2.8 kW. With a raised floor depth of 30 inches (762 mm) and making allowance for necessary subfloor power and data cabling, characteristics of CRAC air plumes, etc., the maximum cooling possible is unlikely to exceed 3 kW per rack unless additional fan-assisted devices are used. In **Figure 5**, we have assumed that five of the 200 racks are high-density racks placed together in a row.

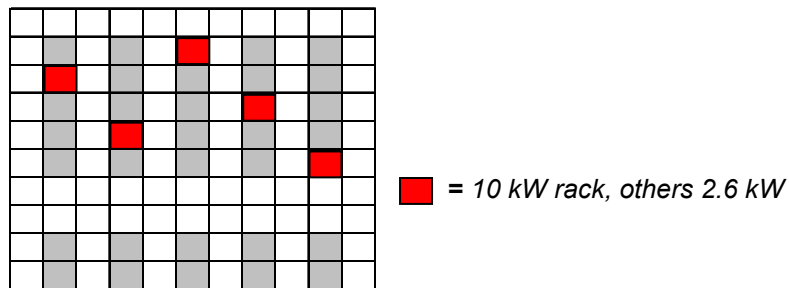
Figure 5 – Data center with all high-density racks together



Assuming that each of the five high-density racks has a load of 10 kW and the remaining 195 have a load of 2.6 kW, the overall average per rack would be 2.8 kW per rack – below the theoretical cooling limit. The average load for the high-density row, however, would be 10 kW per rack, which the cooling infrastructure would be unable to support unless “scavenging” or self-contained solutions were adopted (see later solutions 9 and 10).

An alternative is to spread out high-density racks to maintain the cooling averages as shown in **Figure 6**.

Figure 6 – Data Center with high-density racks spread out

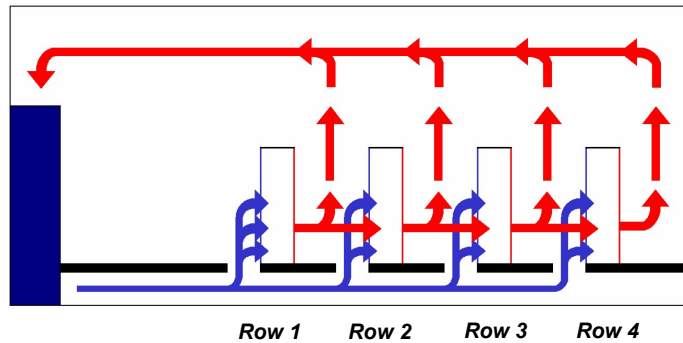


The fundamental reason why spreading out high-density loads is effective is because isolated high power racks can effectively “borrow” underutilized cooling capacity from neighboring racks. However, this effect cannot work if the neighboring racks are already using all the capacity available to them.

6. Implement hot-aisle / cold-aisle arrangement

With few exceptions, most rack-mounted servers are designed to draw air in at the front and exhaust it out at the back. **Figure 7** illustrates a room with all rows facing the same direction. Using this setup, the hot air from row 1 is exhausted into the adjacent aisle, where it mixes with supply or room air and then enters into the front of the racks in row 2. As air passes through each consecutive row the servers are subjected to hotter intake air. If all rows are arranged with intakes facing the same way like this, equipment malfunction is inevitable. This effect occurs in both raised-floor and hard-floor environments.

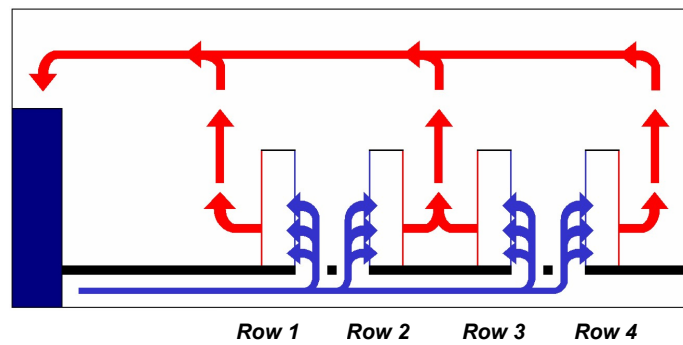
Figure 7 – Rack arrangement with no separation of hot or cold aisles



Best practice is to configure equipment in alternating “hot” and “cold” aisles as shown in **Figure 8**. Cold aisles contain the vented floor tiles, and racks are arranged so that all server fronts (intakes) face a cold aisle. Hot air exhausts into the hot aisle, which contains no vented floor tiles.

This hot / cold aisle arrangement also applies to hard-floor environments. See APC White Paper #55, “Air Distribution Architectures for Mission-Critical Facilities”.

Figure 8 – Hot aisle / cold aisle rack arrangement

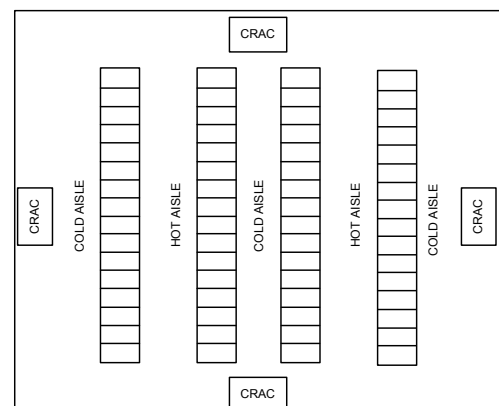


7. Align CRAC units with hot aisles

CRAC units must be aligned with hot aisles to optimize cooling efficiency. **Figure 9** shows a typical room layout where CRAC units have been evenly placed around the perimeter of the room to service a hot aisle/cold aisle arrangement.

In this example, the CRAC units along the two side walls are too close to the cold aisle, which causes the airflow to bypass the floor vents in that aisle. These CRAC units would be better

Figure 9 – Typical CRAC



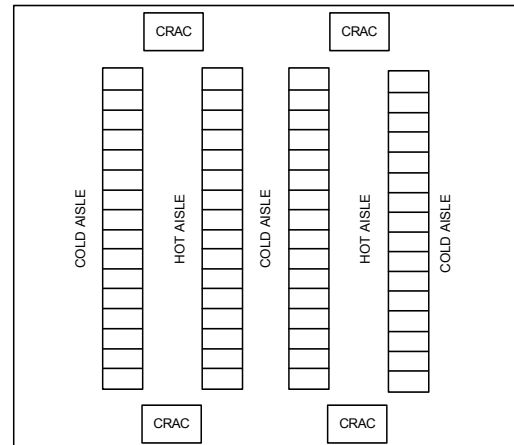
positioned along the top and bottom walls to get better airflow along the aisles.

In **Figure 10** the CRAC units have been moved to the top and bottom walls and are now aligned with the hot aisles.

Conventional wisdom would indicate that CRACs should be aligned with cold aisles to generate a flow of air to the floor vents. However, CFD (computational fluid dynamics) analysis has shown that hot air from the hot aisles crosses over the racks into the cold aisle when returning to the CRACs, causing a mixing of hot and cold air that increases the temperature of supply air to the rack fronts.

In summary, with a raised-floor cooling system it is more important to align CRAC units with the air *return* path (hot aisles) than with the subfloor air *supply* path (cold aisles).

Figure 10 – CRACs aligned with aisles



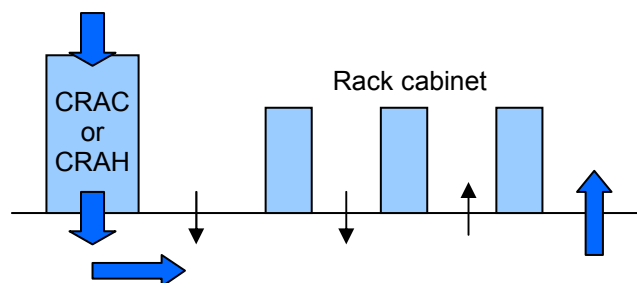
8. Manage floor vents

Rack airflow and rack layout are key elements in maximizing cooling performance. However, improper location of floor vents can cause cooling air to mix with hot exhaust air before reaching the load equipment, giving rise to the cascade of performance problems and costs described earlier. Poorly located delivery or return vents are very common and can negate nearly all the benefits of a hot-aisle/cold-aisle design.

The key to air delivery vents is to place them as closely as possible to equipment intakes, which maximizes keeping cool air in the cold aisles. For subfloor air distribution, this means having vented tiles in cold aisles only. Overhead distribution can be just as effective as subfloor distribution – but again, the key is that distribution vents be located over only *cold* aisles, and for the vents to aim airflow directly downward into the cold aisle (not laterally using a diffusing vent). In either overhead or subfloor systems, any vents located where equipment is not operational should be closed since these sources end up returning air to the CRAC unit at lower temperatures, which increases dehumidification and decreases CRAC performance.

Floor vents placed too close to CRAC units produce negative pressure, causing air from the room to be drawn back under the floor as shown in **Figure 11**. A simple air velocity measuring device can be used to establish vent tile locations that ensure the correct static pressure.

Figure 11 – Relative air movement in high-velocity underfloor environments



Note: Rack orientation may differ in some data centers. The example shown above is different from that illustrated earlier in **Figure 10**, but is included to demonstrate the airflow pattern described above.

The key to air return vents is to place them as closely as possible to equipment exhausts in order to maximize the collection of hot air from the hot aisles. In some cases, an overhead dropped ceiling plenum is used to collect the hot air, which facilitates aligning return vents with hot aisles. When a high, open, bulk return ceiling is used, the best approach is to locate the returns of the CRAC unit as high up in the ceiling as possible and, where possible, spread out the return using ductwork in an attempt to align returns with the hot aisles. Even a crude return plenum with only a few return vents roughly aligned with hot aisles is preferred over a single bulk return at the side of the room.

For smaller rooms without raised floor or ductwork, upflow or downflow CRAC units are often located in a corner or along a wall. In these cases, it can be difficult to align cool air delivery with cold aisles and hot air return with hot aisles. Performance will be compromised in these situations. However, it is possible to improve the performance of these systems as follows:

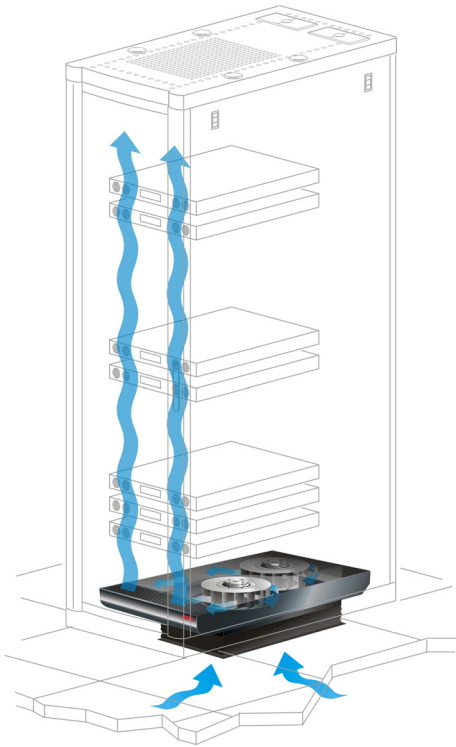
- For upflow units, locate the unit near the end of a hot aisle and add ducts to bring cool air to points over cold aisles as far away from the CRAC unit as possible.
- For downflow units, locate the unit at the end of a cold aisle oriented to blow air down the cold aisle, and add either a dropped-ceiling plenum return or hanging ductwork returns with return vents located over the hot aisles.

A study of poorly placed supply and return points reveals a major underlying root cause: personnel feel that some aisles are hot and some are cold and assume this is an undesirable condition and attempt to remedy it by moving cool air vents to hot aisles, and moving hot air returns to cold aisles. **The very condition that a well-designed data center attempts to achieve – the separation of hot and cool air – is assumed by personnel to be a defect and they take action to mix the air, compromising the performance and increasing the costs of the system. Hot aisles are supposed to be hot.**

9. Install airflow-assisting devices

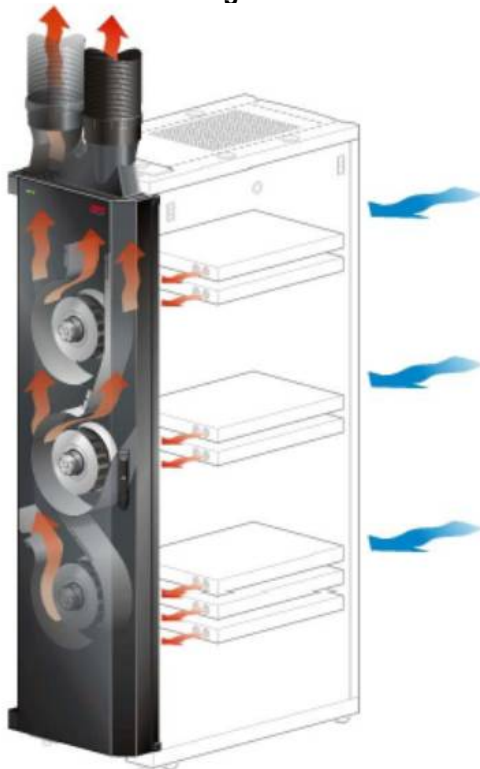
Where the overall average cooling capacity is adequate but hot spots have been created by the use of high-density racks, cooling loads within racks can be improved by the retrofitting of fan-assisted devices that improve airflow and can increase cooling capacity to between 3 kW and 8 kW per rack. Devices such as APC's Air Distribution Unit (ADU) effectively "borrow" the air from adjacent racks (**Figure 12**). As with all air-scavenging devices, care must be taken when positioning the device to ensure that the air taken from the adjacent space does not result in overheating of neighboring racks. These devices should be UPS-powered to avoid thermal shutdown of equipment during power outages. In high-density environments, thermal overload can occur during the time it takes to start the backup generator.

Figure 12 – Rack-mounted fully ducted air supply unit



Fan-tray devices, such as APC's Air Distribution Unit (ADU), fit into the rack's bottom U spaces and direct the airflow vertically to create a cold air "curtain" between the front door and the servers. Blanking panels (see solution #3 earlier in this paper) must be used to ensure the integrity of this newly created plenum.

Figure 13 – Rack-mounted fully ducted air return unit

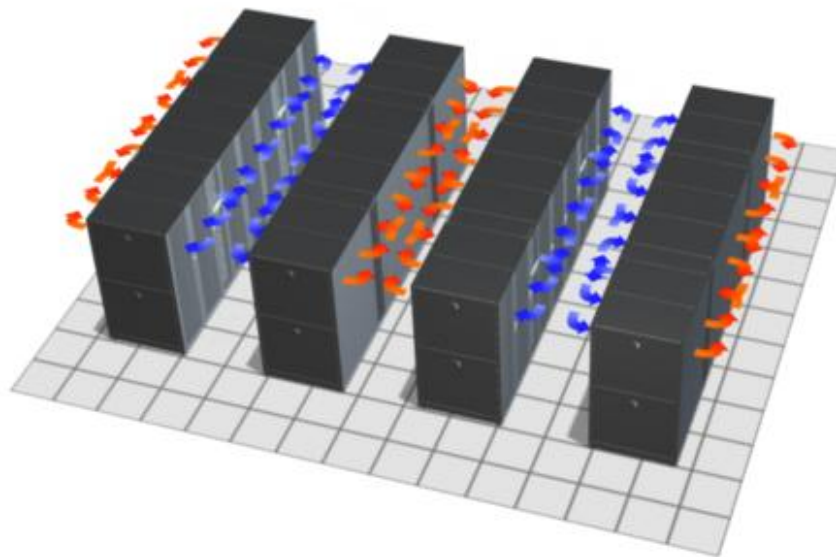


For higher densities, the rear door of the cabinet can be removed and replaced with an air-moving device such as APC's Air Removal Unit (ARU). Hot exhaust air that would normally be expelled into the hot aisle is gathered and propelled upwards, where it is ducted into the return air plenum. This eliminates recirculation at the rack and improves CRAC efficiency and capacity. *Blanking panels and rack side panels must be used with these devices.*

10. Install Row-based Cooling Architecture

As power and cooling requirements within a rack rise, it becomes increasingly difficult to deliver a consistent stream of cool air to the intakes of all the servers when relying on airflow from vented floor tiles (See APC White Paper #46, “Cooling Strategies for Ultra-High Density Racks and Blade Servers” for more information on the limitations of cold air distribution for high density applications). Row-based cooling architecture focuses on heat removal and eliminates the concern of proper cold air distribution from floor tiles. By placing the cooling equipment in the row, the heat is captured and neutralized before mixing in the room. This provides more predictable cooling to IT equipment. APC White Paper #130, “The Advantages of Row and Rack-Oriented Cooling Architectures for Data Centers” discusses the differences and advantages of this approach over traditional data center cooling strategies. **Figure 14** illustrates one example of this row-based cooling architecture.

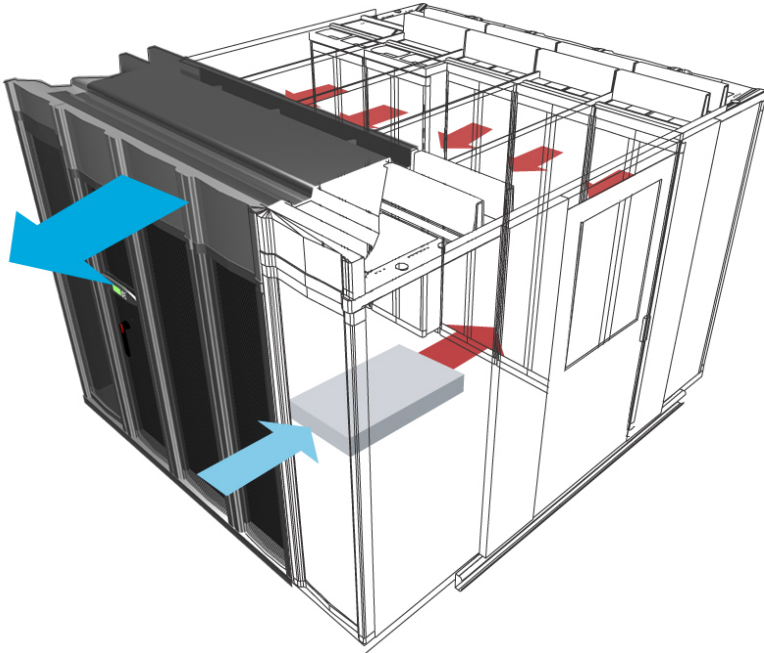
Figure 14 – Row-based Cooling Architecture



To further improve the efficiency and predictability of an row-based system, either rack or row-based containment systems can be added. These self-contained high-density cooling systems are designed to be installed in a data center without impacting any other racks or existing cooling systems. Such systems are thermally “room neutral” and will take hot air from IT Equipment, and discharge cool air back into the room, or enclosed cabinet.

Containment of the hot air from IT equipment increases return air temperatures to the row-based cooling devices, improving the efficiency of the system. This also completely eliminates mixing of hot air, improving the predictability of the row-based system. Two examples of these systems are illustrated in **Figure 15** and **Figure 16**.

Figure 15 – Hot Aisle Containment System (high density zones)



For high density zones, a hot aisle containment system (HACS) can be deployed to contain the hot aisle. Ceiling panels enclose the top of the row, while a set of end doors are used to contain the end of the hot aisle. Hot air from the servers (up to 60 kW per rack) is discharged into the contained hot aisle and drawn through the cooling unit to be discharged back into the room at ambient temperature.

Figure 16 – Rack Air Containment System (supports up to two IT racks)



In a rack air containment system (RACS), single or multiple row-based cooling systems are tightly coupled with the IT enclosure, ensuring the maximum effectiveness of heat removal and cool air delivery to the rack-based equipment. (Up to 60 kW per Rack)

Conclusion

Installation of the latest blade-server technology provides many benefits. However, these servers – if deployed as compactly as their size allows – draw two to five times the per-rack power of traditional servers and generate heat output that can easily cause thermal shutdown if proactive cooling strategies are not employed. To avoid outright equipment failures, unexplained slowdowns, and shortened equipment life, it is becoming critically important to implement a regular health check regime to ensure that cooling equipment is operating within the design values of capacity, efficiency, and redundancy. The solutions outlined in this paper will help keep the data center operating at peak efficiency to maintain the business processes it supports and to prevent future problems.

The guidelines in solutions 1-8 will help keep a typical data center operating within its original design limits. Solution 9 shows how supplemental devices can be used to improve air distribution and, solution 10 discusses a new scaleable cooling architecture that can be implemented to address high density cooling applications.

About the Author:

Peter Hannaford is Director of Product Marketing for APC's EMEA region. He is a member of the British Chartered Management Institute and the British Institute of Directors, and has been involved in the design and construction of over a million square feet of data center facilities worldwide.