



White Paper
Intel Information Technology
Computer Manufacturing
Thermal Management

Air-Cooled High-Performance Data Centers: Case Studies and Best Methods

By combining innovations and best-known methods for air-cooled data center design, we have achieved breakthrough power and heat densities of 15 kilowatts per cabinet and more than 500 watts per square foot of server room area, at a lower cost than other data center designs. Our models indicate that it is possible to double that density in future.

Doug Garday and Daniel Costello, Intel Corporation

November 2006

IT@Intel

Executive Summary

By using a holistic approach to designing an air-cooled data center, we have achieved breakthrough power and heat densities of 15 kilowatts (kW) per cabinet and more than 500 watts per square foot (WPSF) of server room area, at a lower cost than most other data center designs. We believe, based on extensive modeling, that we can double that density in future.

Our results show that it is possible to use air cooling to achieve much greater power densities than we previously thought.

We began by studying air cooling issues and developing design principles based on best-known methods. We progressively refined our approach, incorporating novel techniques as we designed and built successive data centers. These designs included a two-story building, a retrofit of an existing facility, and a data center without an RMF.

Our approach includes several innovations:

- We modeled airflow to identify and address key airflow problems.
- We used barriers and custom cabinets to control airflow and increase the air conditioning airflow efficiency.
- We combined multiple techniques to achieve high densities in new and retrofitted buildings.

Our results show that it is possible to use air cooling to achieve much greater power densities than we previously thought. We believe that these results will further stimulate the debate over whether to use air or water cooling in future data centers.

Contents

Executive Summary	2
Business Challenge	4
Air Cooling Issues	4
Solution	6
Standardized Hot and Cold Aisles	6
Return Air Plenum	6
Barriers	7
Cabling	7
Cabinet Rows Parallel with Airflows	7
Makeup Air Handler	7
Load Characterization	7
Case Study 1: A Two-story Data Center	8
Case Study 2: Blade Servers and Chimney Cabinets	11
Case Study 3: Blade Servers without a Raised Metal Floor	13
Case Study 4: A Future High-Performance Data Center Design	15
Conclusion	17
Authors	19
Acronyms	19

Business Challenge

Intel's data centers support all the company's major computing functions, from engineering to core IT enterprise infrastructure. Over the past five years, the increasing power and cooling demands of servers and other IT equipment have challenged Intel IT to develop strategies for rapidly increasing the power density that our data centers can support.

In designing data centers, we repeatedly face the same questions:

- What is the maximum air cooling capacity that we should build into our data centers based on the anticipated future density and heat characteristics of blade and other servers?
- What are the limits of air cooling and when should we start planning to use an alternative cooling method?
- Can our data centers accommodate the servers that will be produced during the next five to seven years?

The challenges are significant: Blade servers can generate a heat load of 14 kW to 25 kW per cabinet, representing a density of more than 300 WPSF over the total data center area, as defined in the sidebar on page 10. However, many existing and planned data centers support only a fraction of this density; one analysis of 19 data centers estimated average density at the end of 2005 to be only 32 WPSF.¹

When researching the issues, we also discovered that many IT professionals at other companies felt it would be too expensive to upgrade existing data centers to support new generations of servers. Even our own construction department joked that the acronym HPDC (high-performance data center) really stood for "high-priced data center."

¹ "2005–2010 Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment." The Uptime Institute. 2000–2006.

Many in the industry were starting to talk about liquid cooling as the way to solve the server density challenge. We were skeptical.

We were reluctant to give up on air cooling and start down a new path that might ultimately present even more problems. Instead, we chose to study air cooling more closely and look for affordable solutions.

Air Cooling Issues

When we started measuring data center heat densities, we were amazed to discover how inefficient our cooling methods were.

Our computer room air conditioning (CRAC) units were moving huge quantities of air, yet their cooling coils were only providing 10 to 25 percent of their rated cooling capability. There were localized hot spots within the data center and some servers were ineffectively drawing in hot air from the hot aisles. Other areas were very cold, with most cold air bypassing the hot equipment and short circuiting right back to the CRAC units.

In general, we were moving a lot of air and using considerable fan energy, yet we received little benefit from most of the cooling equipment. If cooling air cannot be effectively channeled through hot servers, it is wasted.

For example, we supplied cold air at 55 degrees Fahrenheit (° F) through diffusers in the RMF into the cold aisles between rows of cabinets. However, we found that the air entering the top servers in some cabinets was over 80° F.

This meant that the cold aisle wasn't a cold air delivery zone as intended; it was a zone where servers were drawing in a mix of cold air and the hot air exiting the equipment.

Figure 1 shows this effect using computational fluid dynamics (CFD) modeling of a cross section through a two-story data center with a 65° F cold air supply. It shows how some of the hot air from the hot aisles gets drawn into the cold aisles.

We also developed a related theory that we named the "vena contracta effect," after the term used for the narrowest, fastest part of a jet of liquid or gas. Based on Bernoulli's equations, our theory states that the higher the velocity of the air supplied into a cold aisle, the more surrounding air it will induce into the cooling air plume, as shown in Figure 2 on the following page. Since

the air induced is the hot air exiting the servers, it has a detrimental effect on server cooling.

We considered traditional ways to counter this effect:

- Lowering the supply air temperature, which negatively affected overall cooling efficiency
- Supplying more air to the cold aisle to flush out the hot air, which only induced more hot air into the cooling air plume
- Making the cold aisle wider to slow down the supply air velocity, which lowered the overall server and heat density

All these options were inadequate. To solve these issues and other air cooling challenges, we needed to fundamentally change the way we approached the problem.

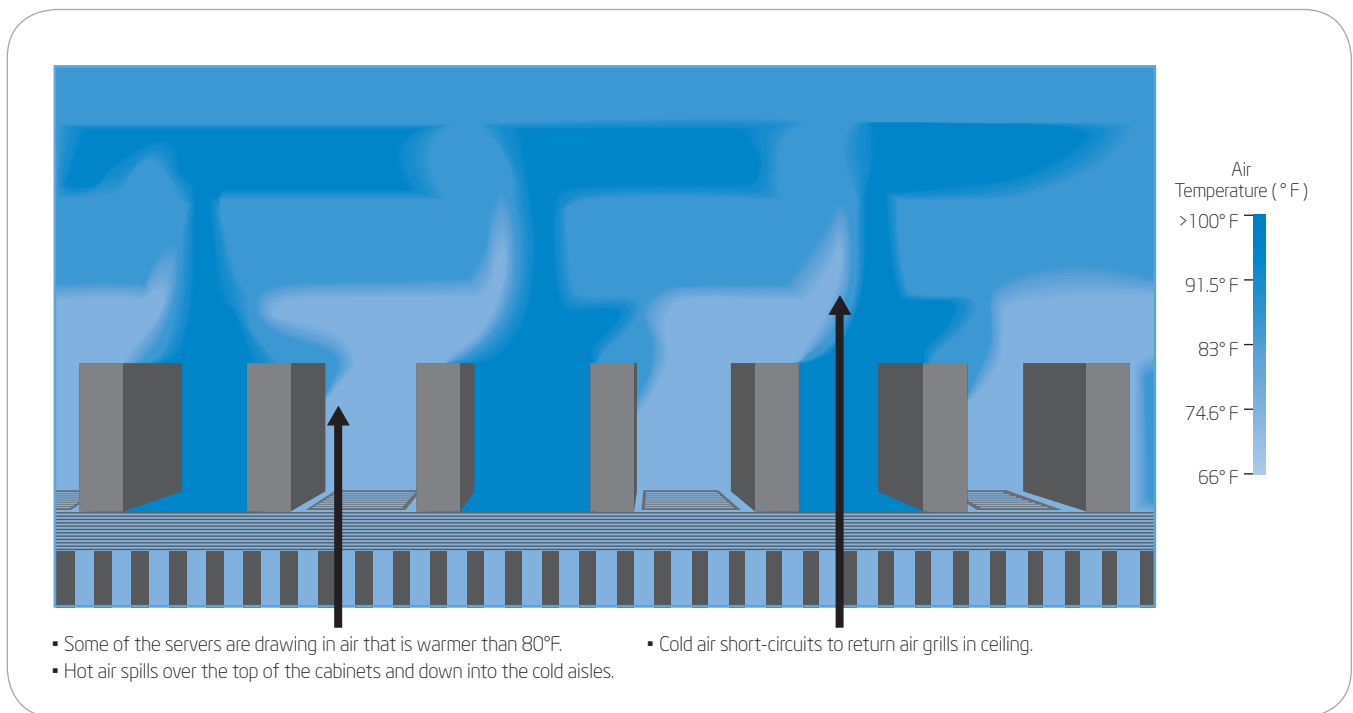



Figure 1. Hot air re-entrainment and cold air bypass in a data center without barriers between the aisles.



Solution

In 2002, we formed a multidisciplinary engineering team to focus on data center density issues. We defined a series of design concepts that positioned us to increase power and heat densities in future data centers. Our holistic approach included barriers between aisles as well as other best practices.

Standardized Hot and Cold Aisles

We always oriented cabinets within rows in a uniform way so that two adjacent rows of cabinets drew their cooling air from a common cold aisle and expelled their hot air into a common hot aisle.

We had to overcome some misconceptions: In some cases, engineers had located floor diffusers in hot aisles, believing they needed to be cooler. This wasted precious airflow and reduced the volume of air delivered into the cold aisles; it also reduced the usable capacity of the CRAC units by lowering return air temperatures. Heat exchange

is based upon the temperature differential between the return air and coil temperature. We emphasized that cold aisles should be cold, hot aisles can be hot, and that we wanted to increase the return air temperature—not lower it.

Return Air Plenum

Typically, data centers are located in buildings with a nominal ceiling height of 14 to 17 feet, below which is a dropped acoustic ceiling located 9 to 10 feet above the RMF. With a 2- to 3-foot RMF, this meant there was a 2- to 6-foot zone above the dropped ceiling that we were not using effectively. We converted this dead air zone into a return air plenum.

- We placed return air grills in the ceiling above the hot aisles to provide a path for the rising hot air to escape into the return plenum.
- We fitted the CRAC units with a duct on top of the air inlet so that they drew hot return air out of the plenum instead of drawing it from the room. This changed the hot air return path: Instead of flowing horizontally across the room, hot air moved vertically up into the return plenum. This reduced hot air re-entrainment.

Barriers

We found that best approach to preventing cold and hot air from mixing was to place barriers above and below the cabinets, isolating the hot aisles from the cold aisles. This radical new technique significantly reduced both the amount of hot air that infiltrated back into the cold aisle (re-entrainment) and the cold air that went directly from the floor diffusers to the return air grills (bypass).

Cabling

The telecommunications cable trays and power wireways sometimes obstructed airflow up through floor diffusers in the cold aisles. We therefore located them under the solid tiles of the hot aisles.

Under-floor cables traditionally entered cabinets through RMF cutouts in the hot aisles. These cutouts were yet another source of leaks because they allowed air to short circuit the equipment. We replaced them with brush-lined openings that allowed the cabling to pass through, yet blocked excess leakage.

Cabinet Rows Parallel with Airflows

In data centers built on an RMF, the diffusers closest to the CRAC units tend to emit the least air or, in some cases, actually draw air from the room. This meant that if we placed rows of cabinets perpendicular to the airflow from the CRAC units,

then the cold aisles closest to the CRACs received the least airflow, while cold aisles in the center of the room received the most. To reduce this effect, we placed rows parallel with the airflow. This also reduced airflow obstructions by orienting the cables under the hot aisles parallel with the airflow.

Makeup Air Handler

We noticed that sometimes one CRAC was cooling while its neighbor was heating. One CRAC was humidifying while another was dehumidifying. There was no coordination between the CRACs and as a result, capacity was wasted. We adopted a coordinated, centralized approach using a makeup air handler (MAH) to manage humidity control for the entire data center. The MAH filtered and conditioned the makeup air, either adding or removing water vapor.

Load Characterization

We carefully studied the loads on the cooling system and characterized the servers and file storage. We measured:

- Airflow through servers
- The difference in temperature between air entering and air exiting the servers (defined as delta-T or ΔT)
- The power that servers used when idle and when under load from test software

We used this data to define our design requirements.

One interesting aspect of this was that servers incorporated fans that switched between low, medium, and high speed depending on the temperature. Since our goals were to minimize airflow and maximize server ΔT , we ensured that we supplied air below these temperature cutoff points so that fans always ran at their lowest speed. This may seem like a trivial point, but we found that it produced a series of benefits, reducing the velocity of airflow and the resulting mixing of hot and cold air.

Case Study 1: A Two-story Data Center

Once we had developed our design principles, we began to apply them to the design and construction of data centers. In 2004, we began designing a data center to contain 220 cabinets fully populated with single height rack mount (1U) servers. Each 1U server required 39.6 standard cubic feet per minute (SCFM) of cooling air at the fan's lowest speed and produced a ΔT of 26.5° F. We planned to fill each cabinet with 40 of these servers, resulting in a total design load of 348,480 SCFM of cooling air.

Our initial design attempt proved to be less than optimal, largely because we departed from our newly developed principles in order to emulate other data centers considered state of the art at that time. The two-story design included a standard hot and cold aisle layout, but lacked heat containment methods such as barriers between aisles.

The problems with this approach became apparent when we conducted CFD modeling during the design process. This made it clear that our initial airflow design would produce many undesirable effects, including a room temperature of more than 100° F. Asymmetrical airflow would cause most cooling air to flow through one side of the room. The relatively high velocity of air entering the cold aisles caused a vena contracta effect, inducing hot air from the server fan outlets to flow beneath, around, and up the face of the cabinets.

To correct these issues, our engineers developed two techniques:

- A symmetrical airflow system that created even pressure differentials everywhere in the room
- Barriers between hot and cold aisles

When we modeled the reconfigured data center, we saw a dramatic improvement. We were now ready to push the envelope in air cooled data centers.

We used our design principles as a basis. When constructing the data center, we plugged cable openings to stop leakage and removed airflow obstructions under the RMF.

Our barrier consisted of a lightweight wall between each hot and cold aisle, extending from the top of each row of cabinets to the drop acoustic ceiling to block airflow across the top of the cabinets, and from floor to ceiling at each end of the aisles. Hot return air flowed through grills in the hot aisles into a return air plenum above the drop ceiling.

We found there were other benefits from this approach. We no longer needed to supply as much air and the air didn't need to be as cold. The air entering the top servers was the same temperature as the air entering the bottom servers, so we had a much more stable cooling system.

We also included other innovations. Our design called for one floor diffuser in the cold aisle

in front of each cabinet. However, standard perforated air diffusers included only a 25 percent free area, which did not allow enough airflow. We took the novel step of using floor grates instead. With a 50 percent free area, these could flow 2,500 SCFM of air each with an acceptable pressure drop.

The average design load for the existing generation of 1U servers was 13.25 kW per cabinet. To cool these would require up to 30 CRAC units with a maximum capacity of 17,000 SCFM each. Rather than adopt this approach,

we used a smaller number of more-efficient recirculation air handlers (RAH) at 37,000 SCFM each, configured for vertical downflow. These included variable frequency drives (VFDs), enabling us to match capacity to load and providing spare capacity. We could run all the units at a reduced flow and if one should fail or need service, the other units could ramp up to pick up the load.

Even though our goal was to eliminate bypass completely, we knew there would be some

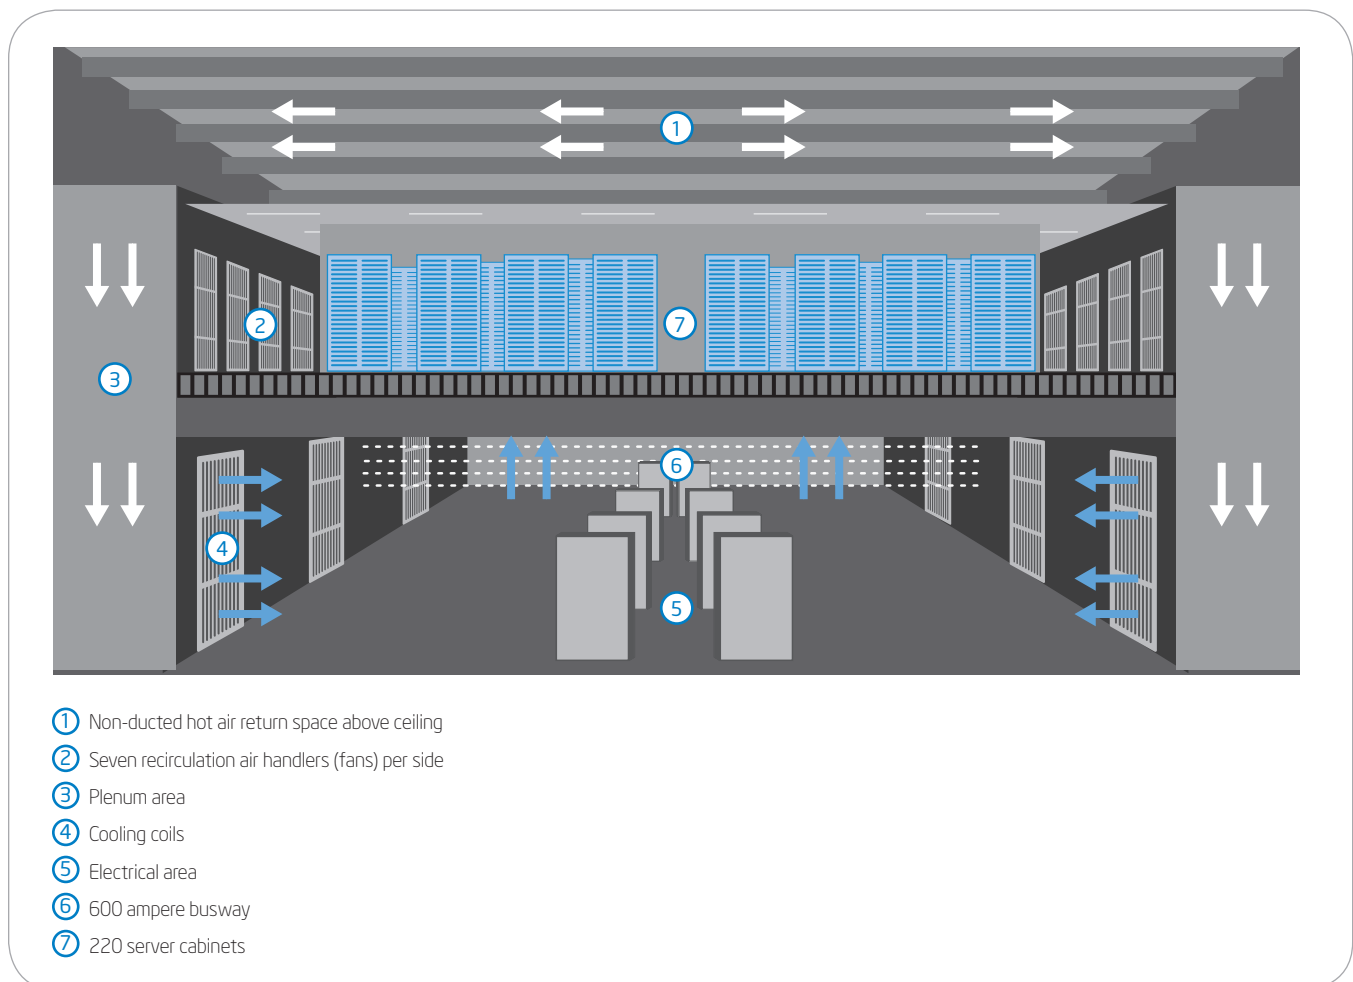


Figure 3. Cross section of a two-story high-performance data center.

leakage, so we designed for 38 percent leakage when using 1U servers.

We used a wet side economizer system, enabling us to exploit evaporative cooling and satisfying a local building code requirement for energy conservation. To make this system more cost effective, we needed to be able to use chilled water at a relatively high temperature. The efficiency of our design enabled us to use supply air at a relatively warm 68° F, which in turn allowed us to use chilled water at 55° F. We could produce water at this temperature with

a wet side economizer system during late fall through spring, which made it cost effective.

This data center represented a breakthrough when it was commissioned in early 2006. We achieved a power density averaging 13,250 watts per cabinet or 496 WPSF over the total RMF area, as shown in Table 1 on page 18. This power density was several times greater than the density goals of many current data center designs. Yet we had achieved this, in a data center delivering Tier 2 reliability and capable of Tier 3, for less than USD 9,000 per kW—well

How to Calculate Power Density

There's a lack of industry standards for calculating data center power and heat density. This makes it hard to obtain consistent measures and comparisons of WPSF. For example, two data centers may have identical equipment, but if one occupies a 1,000-square-foot room and the other a 1,300-square-foot room, the first will have a higher WPSF. There are many other variables. Air handlers can be located directly beside the data center rather than on the RMF; to calculate WPSF, should you include only the RMF area, or do you include the area for the air handlers, too?

We developed three measures that we use to accurately and consistently compare WPSF:

- **WPSF within a work cell.** We have defined this standard measure to normalize power and heat density on a per-cabinet basis. A work cell is the area dedicated to one cabinet.

For example, if a cabinet is 2 feet wide and 4 feet deep, it occupies 8 square feet. The cold aisle in front is 4 feet, shared between two rows of cabinets, so the portion dedicated to each cabinet is 8 divided by 2, or 4 square feet. The same calculation applies to the hot aisle behind each cabinet. In this example, the work cell dedicated to each cabinet is 16 square feet.

Several of our data centers are laid out this way, with an 8-foot distance, or pitch, between the center of the cold aisle and the center of the hot aisle to accommodate cabinets up to 4 feet deep. Another advantage of the 8-foot pitch is that it allows us to coordinate the drop ceiling grid with the floor grid so that standard 2 by 4 foot light fixtures and fire sprinklers can be centered above the hot and cold aisles.

- **WPSF over the RMF/server room area.** We define the RMF area (or server room area in data centers lacking RMF) as the total area occupied by work cells plus the perimeter egress aisles.
- **WPSF over the total data center area.** To obtain a true measure of the efficiency of data center layout and design, we include the area occupied by air handlers, whether or not they are located on the RMF. Therefore, when calculating WPSF for the total data center area, we include the CRAC/RAH area and egress aisles, but not exterior ramps or staging areas.

Our rule of thumb for an efficient layout is 25 cabinets per 1,000 square feet of RMF and air handling equipment area. Therefore, for this paper, we normalized the layout based on 25 to 26 cabinets per 1,000 square feet, which includes the area for cabinets, aisles, networking, power distribution, and air handling equipment.

below industry benchmarks for Tier 2 and less than half the typical Tier 3 cost.²

The hot aisle enclosures helped us achieve a very high level of air conditioning airflow efficiency (ACAE), defined as the amount of heat removed ($\text{watts}_{\text{heat}}$) per standard cubic foot of airflow per minute, also shown in Table 1. We also achieved excellent air conditioning effectiveness, as measured by the ratio of air conditioning power

to computer power. With the economizer, this ratio fell to 0.23, a figure considerably lower (lower is better) than published data from 13 other data centers reported in a benchmarking study.³

Overall, the data center occupied 70,000 square feet spread over two stories, with five 5,900-square-foot modules of RMF area and approximately 3 megawatts (MW) of usable power per module.

² "Tier Classifications Define Site Infrastructure Performance." The Uptime Institute. 2006.

³ "Data Centers and Energy Use—Let's Look at the Data." Lawrence Berkeley National Laboratory.

Case Study 2: Blade Servers and Chimney Cabinets

For this project, we retrofitted an existing 65 WPSF data center to handle high-density servers. The constraints imposed by adapting an existing facility presented several challenges:

- An existing 18-inch RMF with limited floor to deck height
- The need to use an existing chilled water system
- An under-floor chilled water and telecommunications tray that restricted air movement
- 2,100 square feet of landlocked space
- The need to reuse three existing CRAC units to save money
- A perimeter wall with available space for five new 15,200 CFM CRAC units

When we began this project, our design focus was on supporting 1U servers. As in Case Study 1, our goal was to control bypass and eliminate re-entrainment completely in order to increase the ΔT of our air conditioning system and utilize more of our installed coil capacity.

However, during 2005 we realized that we faced a whole new set of design challenges due to Intel's adoption of a new type of server, the blade server. When we characterized blade servers, we found that they generated a 60° F ΔT under peak load—more than double the 26.5° F of 1U servers.

To handle anticipated increases in density, we had developed a new type of cabinet, which we used for the first time on this project. Called the passive chimney cabinet, it performed the same functions as a hot aisle enclosure, but for a single cabinet. The chimney cabinet has a standard open front, but the rear is enclosed by a solid hinged door. A large opening in the top of the cabinet, toward the rear, allows hot air to exit. A short section of ductwork connects this opening to the return air plenum, eliminating the possibility of hot air re-entrainment and controlling bypass, as shown in Figure 4 on the next page.

The air pressure in the return air plenum is slightly lower than the pressure in the room, which helps draw the air out of the cabinet. We found that an additional major benefit was that the flue effect created by hot air rising up the chimney created a strong draft that drew in extra air from the room through the cracks between the blade chassis.

When we characterized six chassis containing a total of 84 blades in one cabinet under load, we measured a peak power consumption of 21.3 kW, or 3.55 kW per chassis of 14 blades each. We also discovered that the 60° F ΔT across the blades was reduced to 50° F when we measured the entire flow path of the passive chimney cabinet. This meant we were drawing in 20 percent of the total airflow through gaps between blade chassis.

We could have plugged those gaps, but decided there was a greater benefit in leaving them open. Because we were using off-the-shelf CRAC units capable of cooling air by 25° to 29° F, we deliberately designed the heating, ventilation, and air conditioning (HVAC) system to bypass

room air and mix it with the hot air from the blade server cabinets. This cooled the return air to a temperature at which the CRAC unit coils could operate.

Our data center included several other key features:

- To limit under-floor velocity and pressure, we supplied some of the cooling air above the RMF by extending the stands supporting the CRACs to raise them above RMF height.
- The supply air temperature was adjustable from 60° to 68° F to offset radiant heat emanating from the skin of the cabinets. We didn't want the room to be too cold. This could impact ergonomics and negatively affect the personnel installing the cabinets.
- The room is extremely quiet for the performance obtained. The rear of each cabinet, where the server fans are located, is enclosed in the flue so that some of the fan noise is contained within the cabinet.




Figure 4. Passive chimney cabinets.

- We achieved an average density of close to 13 kW per cabinet, with a maximum of 15 kW.

When commissioned in 2005, this project was considered very successful for a number of reasons.

- In terms of power density per unit volume of building space, it was our highest performance data center.
- It showed that passive chimney cabinets effectively eliminated re-entrainment of hot air. This meant that we didn't need to blow a column of cold air up the front of the cabinet to flush out hot air; cabinets could draw cooling air from other areas within the room.
- It demonstrated that we could use an overhead busway electrical distribution system to deliver high power density.

- It reduced data center noise because the server fans were contained in a partially enclosed cabinet.

- It proved that we could manage bypass and use off-the-shelf CRAC units with standard 27° F sensible cooling coils for IT equipment producing a 60° F ΔT . We were able to load the coils to their full capacity.

Though the room's dimensions were not ideal for layout efficiency and the hot aisles were 5 feet wide, a foot wider than usual, we were still able to accommodate a respectable 25.2 cabinets per 1,000 square feet. After proving the passive chimney cabinet as a means to manage airflow, we were now ready to explore new air delivery methods and control bypass air even further.

Case Study 3: Blade Servers without a Raised Metal Floor

Following our success with supplying some air above the RMF, we supplied all the cooling air this way in our next major design in 2006. Our design eliminated the RMF and ramps, placing cabinets directly on the slab and providing all cooling air from ductwork and diffusers above the cold aisles. We designed this ductwork to function like a RMF, sharing airflow from multiple RAHs in a common supply air plenum and capable of continuing to operate in the event of a single RAH failure. We placed high-capacity RAHs on a mezzanine, allowing them to draw the hot return air from above and supply air at the level of the dropped ceiling. With no RMF, we supplied telecom and power using dual overhead busways.

We initially included 120 computing cabinets in the design, each with an average load of 15 kW and maximum of 22 kW per cabinet. We designed the data center to be able to expand

to 222 cabinets in the future. The power supply was 1.8 MW per module, with the ability to increase that later.

The project faced many other design challenges, including the need to exploit existing facilities to the best advantage.

- We used a mixture of chimney cabinets and hot aisle enclosures, which enabled us to reuse existing open-backed cabinets that lacked chimneys.
- We designed a new chilled water system to support the data center modules and other areas of the building, and retrofitted the existing chilled-water system components into a de-coupled wet side economizer system to provide year-round partial sensible cooling. We achieved this by separating the condenser water systems to allow the new system to operate at the optimum chiller condenser water temperature. We configured the existing system to produce lower temperature condenser water that was then pumped through a second coil in the RAHs. This pre-cooled the return air, reducing the chilled water load.
- The RAH unit design included a 36.1° F ΔT chilled water coil and a capacity of 50,000 actual CFM (ACFM) with VFDs. All RAHs normally operate at reduced speed and airflow, but if one RAH fails or requires service, the others can ramp up to compensate. The coils provide a relatively high ΔT with low airflow, saving fan energy and reducing noise. This also provides operational flexibility to fine-tune the HVAC system for future generations of servers.
- The choice of coils and HVAC system ΔT represent our most aggressive HVAC design to date. Rewards include reduced RAH size and cost, warmer return air, better heat transfer, better overall chilled water system efficiency, and lower operating costs.

- We supplied the cooling air above each cold aisle using acoustically lined twin supply ducts and a continuous supply air diffuser. We spaced the ducts a foot apart, using the space between them for fire sprinklers, temperature sensors, and light fixtures. This increased the cold aisle width from 4 to 5 feet.
- We also increased the hot aisles to 5 feet, including a central strip containing fire sprinklers and lighting. The 2-foot-wide areas on either side of the strip included either egg-crate return air grills for the hot aisle enclosures or solid ceiling tiles with chimney flue ducts where we installed passive chimney cabinets. We placed barometric return air dampers above the egg-crate return grills. We can adjust these with the RAH flow to bypass more or less room air, tempering the return air to meet the RAH coil capability, yet still keeping the data center at a positive air pressure relative to surrounding areas.

Modeling showed that the design works well in a building 25 or more feet tall, which can easily accommodate the mezzanine with the RAHs. The higher cooling system ΔT also helps defray construction costs by reducing the size of the RAH and ductwork required.

However, eliminating the RMF did not prove to be the engineering coup that we had initially hoped would reduce costs. The supply ductwork was more difficult to design than we first thought it would be and it increased roof loading because it was supported from above. We are still debating the no-RMF concept and need to do more analysis to compare true costs with using an RMF. We feel that an even better way to eliminate ramps may be to depress the slab and use an RMF to deliver supply air.

Case Study 4: A Future High-Performance Data Center Design

In our most recent design, which we have yet to build, we set a goal of using our experience to design an air-cooled data center with even greater WPSF while meeting specific design parameters.

- With the higher ΔT of blade servers in mind, we aimed to increase overall air conditioning system ΔT over previous designs and increase airflow efficiency.
- We wanted to design a data center that would fit into an office building with a ceiling height of only 17 feet.
- We also wanted to maximize layout efficiency for 220 cabinets yet still allow ample room around the rows of cabinets for moving equipment in and out.

We refined our initial design through modeling iterations to provide up to 30 kW per cabinet—double the density achieved by our previous designs—and 1,250 WPSF over the RMF area, at which point our main concerns were total airflow and noise.

Our design is based on a 60- by 88-foot RMF area, as shown in Figure 5. On either side of this is an 88-foot-long air handling unit (AHU) room containing five large RAH units with VFDs operating at 48,275 SCFM each and able to ramp up to 54,000 SCFM to enable continuous operation during preventative maintenance or in the event of a single component failure.

Each cabinet needs 1,579 SCFM to pass through the blades at a 60° F ΔT . Including the air that

passes between the blade chassis, the total rises to 1,894 SCFM. This induction and mixing reduces ΔT from 60° F to 50° F. The result is a total flow for the 220 cabinets per room of 416,680 SCFM.

We plan to use passive chimney cabinets, which work well with warmer supply air temperatures. Warming the 65° F supply air by 50° F as it passes through the cabinets will result in return air at 115° F. We may decide to mix this with an additional 19 percent bypass air (66,065 SCFM), tempering the return air to 108° F to match the capabilities of the RAH coils.

We also expect to include features from our previous designs, such as a makeup air handler, a symmetrical layout, and high-flow floor diffusers with approximately 50 percent free area. All air leakage into the room through the RMF contributes to cooling and planned bypass.

The walls separating the RMF server area from the AHU rooms will extend from the RMF to the structure and we will build the RAHs in a vertical downflow style, sized to return air through this wall. The vertical height available allows space for airflow silencers on both sides of the fan if required. Another option is to use a continuous array of fans along the AHU wall to reduce low-frequency noise levels without

the need for silencers. By drawing the return air through openings in this wall, the AHU room could be maintained at about 75° F instead of 108° F, and the local control panel and VFD electrical components would not require supplemental cooling.

Our modeling has resulted in some other key findings:

- We determined that a 36-inch-high RMF works best, based on modeling of maximum under-floor static pressure and varying airflow from the diffusers at different locations. An RMF lower than this required too much static pressure from the fans; a higher RMF provided minimal benefit for the extra cost.
- Based on this, we could fit the data center into an office building with a 17-foot height from concrete slab to the floor above. This

would accommodate a 3-foot RMF, 10 feet from the top of the RMF to the drop ceiling, and 4 feet for the return air plenum.

- We could design a data center into a new building with 14-foot ceiling height by depressing the slab 3 feet. This would eliminate any need for ramps. The AHU rooms would be further depressed to a total of 42 inches. This allows better airflow from the RAH into the space below the RMF. Since the chilled water piping is also in the AHU rooms, it provides a well for water containment in case of a leak.

This design summarizes our holistic approach. No single factor is responsible for allowing us to double power density. Instead, we expect to achieve this dramatic increase by assimilating the many lessons we have learned over the past five years.




Figure 5. Model for a data center with 30 kilowatt cabinets.

- The design includes excellent layout efficiency, allowing more compute power per unit area.
- Improved airflow management allows us to supply air at 65° to 70° F instead of the traditional 50° to 55° F, yet improve temperature distribution and cooling to the servers. This creates energy savings.
- Higher supply air temperatures result in the air handlers performing only sensible cooling—rather than also performing latent cooling that dehumidifies the data center, creating a relative humidity that is too low for some types of IT equipment.
- Raising the supply air temperature also raises the return air temperature, allowing us to get more cooling capacity from each coil at the same chilled water temperature.
- Increasing the airflow per square foot also increases cooling. Larger capacity RAHs designed in a vertical arrangement provide more airflow and better utilize space.
- The higher ΔT of blade servers means more heat transfer per CFM of air, which reduces the fan horsepower required.
- Eliminating wasted cold air bypass allows us to use that portion of the total airflow to cool equipment.

As an added benefit, the passive chimney cabinets reduce overall data center noise levels.

Conclusion

Over the past five years, we have developed a holistic design approach to achieve breakthrough power densities in air-cooled data centers. Our current designs support 15 kW per cabinet, or about 500 WPSF over the total RMF area.

We achieved these advances by incorporating the many techniques we discovered for improving air cooling efficiency. Even though each of these techniques provides a relatively small improvement in performance, their combined effect is dramatic. Based on extensive modeling, we believe that we can double the density of our current data centers in future designs. Our experience to date is that our models are accurate in predicting data center capabilities.

Our results show that it is possible to use air cooling to support much higher densities of current and anticipated server technology at relatively low cost per kW of IT equipment load. We expect that these results will stimulate further the debate over whether water or air cooling is more appropriate for high-performance data centers.

Table 1. Characteristics of the data centers in our case studies.

	Case Study 1	Case Study 2	Case Study 3	Case Study 4
Raised Floor Height	18 inches ⁴	18 inches	None	36 inches
Layout Efficiency (Cabinets per 1,000 Square Feet (SF))	23.4	25.2	22.2	26.0
Number of Cabinets	220	53	120 (Phase 1)	220
Average Watts per Cabinet	13,250	12,868	15,000	30,000
Maximum Design Watts per Cabinet	17,000	15,000	22,000	40,000
Total Server Power and Cooling Capacity	2,915 kW	682 kW	1,800 kW	6,600 kW
Total Room Area (Including Recirculation Air Handler)	9,408 SF	2,100 SF	5,408 SF	8,448 SF
Watts per Square Feet (WPSF) over Total Area	310	325	333	781
Server Room Area (Typically with a Raised Metal Floor [RMF])	5,880 SF	1,800 SF	3,536 SF	5,280 SF
WPSF over RMF/Server Room Area	496	378	509	1,250
Work Cell Size	16 SF	18 SF	18 SF	16 SF
Total Work Cell Area	3,520 SF	954 SF	2,160 SF	3,520 SF
WPSF over Work Cell Area	828	715	833	1,875
Number/Capacity of Air Handling Units (includes N+1 unit)	14 at 37,000 actual cubic feet per minute (ACFM) each	3 at 12,400 ACFM each 5 at 15,200 ACFM each	6 at 50,000 ACFM each	10 at 54,000 ACFM each
Number/Capacity of Cooling Air Diffusers	220 at 2,186 CFM each	46 at 1,719 CFM each	120 at 1,389 CFM each	240 at 2,011 CFM each
Maximum Delta-T of IT Equipment	60° F ⁵	60° F	60° F	60° F
Delta-T of Air Handling System	20° F	27.3° F	34° F	43° F
Air Conditioning Airflow Efficiency ($W_{heat}/SCFM$)	6.4	8.6	10.8	13.7
Ratio of Uninterruptible Power Supply (UPS) Output to HVAC Power (kW of HVAC/kW of UPS Output)	0.35 (0.23 when running on wet side economizers)	0.33	0.31 (0.20 when running on wet side economizers)	0.30

⁴ Above second floor of two-story design; 20 feet above first floor.

⁵ Originally designed for 1U servers with a delta-T of 26.5° F; now blade servers are being installed.

Authors

Doug Garday is a data center design engineer with Intel Information Technology.

Daniel Costello is the manager of data center designs and standards with Intel Information Technology.

Acronyms

ACAE air conditioning airflow efficiency

ACFM actual cubic feet per minute

AHU air handling unit

CFM cubic feet per minute

CFD computational fluid dynamics

CRAC computer room air conditioning

ΔT delta-T

HPDC high-performance data center

HVAC heating, ventilation, and air conditioning

MAH makeup air handler

RAH recirculation air handler


RMF raised metal floor

SCFM standard cubic feet per minute

UPS uninterruptible power supply

VFD variable frequency drive

WPSF watts per square foot



www.intel.com/IT

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel, the Intel logo, Intel. Leap ahead. and Intel. Leap ahead. logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2006, Intel Corporation. All rights reserved.

Printed in USA
1106/ARM/RDA/PDF

Please Recycle

Order Number: 315513-001US