hp industry standard servers

october 2002

technology brief TC021002TB

Power and Cooling Trends in the Datacenter

table of contents

abstract

Customers who are building a new datacenter facility or upgrading an existing facility must be aware of the constantly changing power and cooling requirements of computer hardware. This paper provides HP's best estimates of future power and cooling requirements for some HP datacenter products and it describes methods to maximize cooling effectiveness in high-density datacenters. The actual power and heat densities that customers should plan for will depend on the equipment specification and the customer's IT strategy and hardware adoption rates.

introduction

As a result of rapid technology advances, datacenters have evolved from pristine air-conditioned rooms housing a few mainframe computers to high-density computing centers with hundreds of servers requiring the power and cooling capacity of a small town. For organizations that are building or upgrading a datacenter, the challenge is to plan and design the datacenter without knowing which technologies it will house five to ten years from now.

This technology brief describes trends affecting datacenter design, explains how to determine power and cooling needs, and describes methods for cost-effective cooling. The first section explains some basics of electrical power that are necessary to understand the power consumption of computing devices.

power and power factor

Electricity is generated and distributed by utilities as three-phase alternating current (AC), usually at a voltage in excess of 150,000 volts (V). By the time the current reaches the datacenter, its voltage has been stepped down to 600 V or less. At the facility, transformers convert the three-phase current into three separate phases. The voltage of each phase is represented by a sinusoidal wave that alternates between positive and negative values (see Figure 1) at a frequency of 60 cycles per second (or 50 cps in many European countries). The current in each phase has a similar waveform at the same frequency. Most computer equipment operates on single-phase current. Single-phase loads, such as computer equipment, are connected to one of the transformer's phase windings and the Neutral connection. Equipment requiring more power, including datacenter environmental support systems, runs on three-phase current. Three-phase loads, such as air-conditioning equipment, are connected to all three transformer windings.

figure 1. Representation of sinusoidal voltage waves of three-phase current and corresponding transformer phase windings

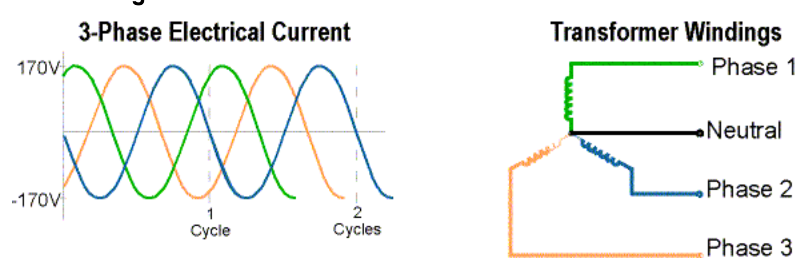

Before computing and storage devices can use electrical power, the alternating current must be transformed to direct current by a power supply. The term *power* is the rate at which the electricity does work, such as running a CPU or turning a cooling fan. The power that the electricity provides (Apparent Power) is simply the voltage times the current, measured in volt-amperes:

Apparent Power = voltage x current [VA]

There is a difference between the power supplied to a device and the power actually used by the device because of the capacitive nature at the input of the device to delay current flow. This delay causes the current to lag behind the voltage and the current to be narrow and high pulses (Figure 2). The difference between the two wave peaks is called the phase angle, which is measured in degrees; one cycle is equal to 360 degrees.

figure 2. Phase angle is the difference between the voltage and current wave peaks.

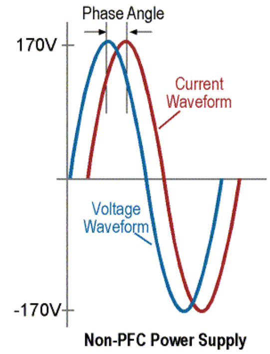

The true power used by the load is measured in Watts (W). The relationship between true power and apparent power (AP) is as follows:

Power [W] = cosine(phase angle) × Apparent Power [VA]

Cosine(phase angle) is known as the Power Factor of the load. Power Factor is a number between 0 and 1 that indicates the timing relationship between the voltage and current, in other words, how efficiently the current is converted to useful work. A power supply that has a phase angle of zero (the voltage and current peak together) has a PF of 1.0, which results in the most efficient loading of the device.

Power supplies for servers usually contain circuitry to "correct" the Power Factor (making input current and voltage to be in phase). Power Factor correction allows the input current to continuously flow, reduces the peak input current, and reduces the energy loss in the power supply, thus improving its operation.

Power factor-corrected (PFC) power supplies have a power factor near unity (~ 1), and thus are more efficient (see Figure 3). The use of energy-efficient PFC devices, including uninterruptible power supplies, can lead to significant cost savings.

figure 3. Voltage and current waveforms of a power factor-corrected power supply

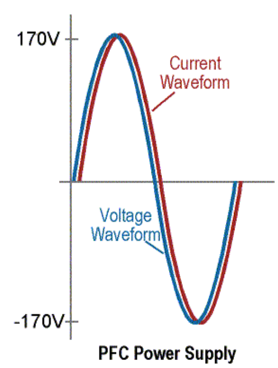

power and cooling trends

When datacenters mainly housed large mainframe computers, power and cooling design criteria were designated in average wattage per square foot (or square meter) and British Thermal Units per hour (BTU/hr). These design criteria assumed that power and cooling requirements were equal across the entire datacenter. Today, datacenters are populated by racks of scalable computing systems that require enormous amounts of electricity and produce unimaginable amounts of heat. Average (per unit area) design criteria do not encompass the specific power and cooling requirements of high-density solutions. Consequently, datacenters can no longer be designed by using average wattage and BTU criteria.

The trends toward denser, compute-intensive systems will continue. Therefore, customers who are designing new facilities or retrofitting existing ones need to be aware of these trends to optimize their infrastructure investment. The following sections describe some of the trends in power distribution and cooling requirements.

modular power distribution units

In traditional datacenters, power goes from the transformers to a number of sub-panels that contain circuit breakers. A wire from each circuit breaker provides power to each outlet where the equipment is connected. As long as the power provided is sufficient for the equipment in place, this system is adequate. However, if the outlets, wire, and breakers need to be upgraded for higher amperage or three-phase power, the process can be very expensive and time-consuming. A device called a modular power distribution unit (PDU) can alleviate this problem (Figure 4). Modular PDUs integrate the outlets, wire, and breakers in a convenient location on each rack of equipment.

figure 4. Modular PDU in zero-U configuration

HP modular PDUs range from 24 to 40A and have up to 32 outlet receptacles. For mission-critical environments with redundant power systems, HP also offers a fault-tolerant Dual Input PDU that automatically switches over to a secondary input source if the first source fails. These PDUs can be mounted in a zero-U or 1U configuration without consuming valuable rack-space (1U = 1.75 inches, 4.45 cm).

uninterruptible power supplies

Datacenters cannot rely on utility power grids as a source of continuous power for critical equipment. For applications that require a high level of fault tolerance, the power source for the equipment is protected by an uninterruptible power supply (UPS). The UPS is connected between the equipment and the electrical outlet or PDU. There are different types of UPSs, but all provide a primary and secondary power source for the equipment. The UPS contains batteries that protect against power disturbances—natural or man-made—and circuitry that conditions (filters and enhances) the utility power to provide more stable current. The capacity of a UPS is mainly determined by the size of its battery. The larger the battery, the longer the connected equipment can operate before shutting down.

To meet the 24x7 requirements of the datacenter, HP makes modular UPSs that can be scaled to meet the demand for more power. Additional battery modules can be added, or optional UPS card slots can be used, to increase the capacity and backup time. Through the use of a UPS card slot, additional UPSs to be connected to the original UPS to scale capacity. For high-density datacenter environments, HP rack-mountable UPSs offer remarkable power density (Watts/U-space) for more capacity in a space-saving form factor.

These intelligent UPSs are bundled with HP Power Management Software to extend the operation of the mission-critical devices by using load segment control to schedule startups and shutdowns of less critical devices.

racks

For years, computing equipment has been stacked in vertical racks to reduce datacenter floor space requirements, but rack mounting equipment has had the opposite effect on power consumption and heat density. Figure 5 shows historic heat density data and future trends for fully-configured telecommunications and data processing racks from 1992 to the year 2010. The averages shown are based on the heat dissipated by the equipment, divided by the footprint of the equipment. The values must be adapted to the average for the total space. These trends are based on reasonable estimates and should not be construed as being accurate.

figure 5. Product heat density trends chart

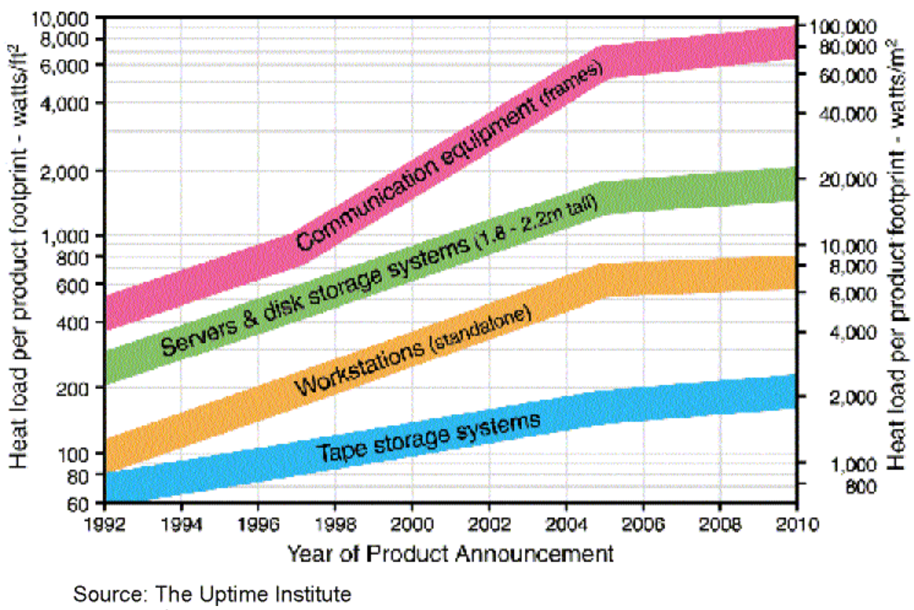

Source: The Uptime Institute

Racks range from 70 to 87 inches (1.8 to 2.2 meters) in height, which means they can hold from six 8U servers to forty-two 1U servers, not including storage or UPSs. A typical 42U cabinet with dual processor (2P) servers and storage requires over 12 kW of power (see Table 1). The heat load for this rack is almost 42,000 BTU/hr (about the same as for a typical one-story house); however, the cooling requirement for a datacenter full of such racks is enormous. Now consider that IA-64 servers and higher capacity drives may increase the power requirement to 15 kW per rack.

In high-density deployments, perforated rack doors are required to ensure that internal cooling fans can pull sufficient air into rack-mounted servers. Racks housing ProLiant servers must have perforated doors with a minimum open area of 63 percent. Blanking panels should also be used to close off empty areas of a rack to prevent hot exhausts from circulating inside the rack.

A growing trend for building new IT infrastructures involves ultra-thin 1P blade servers that can be stacked in a chassis, interconnected, and easily managed (see Table 1). ProLiant BL 10e servers were the first to enter the market. This technology promises to lower the operating cost per CPU by reducing management expenses and the requirements for floor space, power, and cooling.

table 1. Power and cooling requirements* of fully-configured, density-optimized ProLiant servers

	DL 360	DL 380	DL 580	BL 10e
ProLiant Server	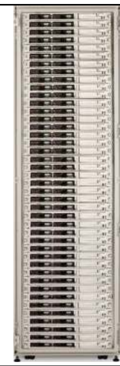			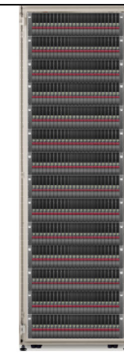
Servers per Rack	42	21	10	280
No. of CPUs	2P × 42 = 84P	2P × 21 = 42P	4P × 10 = 40P	1P × 280 = 280P
Power	42 × 292W = 12.26kW	21 × 200W = 4.2 kW	10 × 760W = 7.6 kW	14 × 700W = 9.8 kW
Heat Load**	41,820 BTU/hr	28,700 BTU/hr	25,956 BTU/hr	33,440 BTU/hr
Cooling Requirement***	3.5 tons	2.4 tons	2.2 tons	2.8 tons

* These calculations are based on the product nameplate values for fully configured racks, which may be higher than the actual power and cooling requirements.

** To determine the heat load (BTU/hr) of a particular device, multiply the watts it consumes by 3.413 BTU/hr.

*** In the U.S., the cooling requirement is expressed in "tons"; a standard ton of refrigeration corresponds to a heat absorption rate of 12,000 BTU/hr.

power supplies

Power supplies are one of the most critical components in computing and storage devices. Power supplies convert AC power into DC power at various voltage levels (12 V, 5 V, 3.3 V, -5 V, and -12 V) based on the requirements of internal components. Power supplies also conserve power by reducing the amount of power consumed by components when they are idle. A power supply not only provides power to components, but it is also their main source of cooling. It contains the main fan that controls the flow of air through the equipment. For this reason, the power and cooling capacity of the power supply determines the number of drives and processors that can be used.

power supply capacity is growing

The capacity of a power supply is the total output power of all the voltages that it provides. The total output power is calculated by adding the products of the voltage and current for each voltage level. The capacity of power supplies is increasing to accommodate the increased power consumption and heat generation of components such as drives and processors (see side bar). In 1992, the ProLiant 4000 server used a 540 W power supply. Today, ProLiant BL servers use 3000-W power supplies.

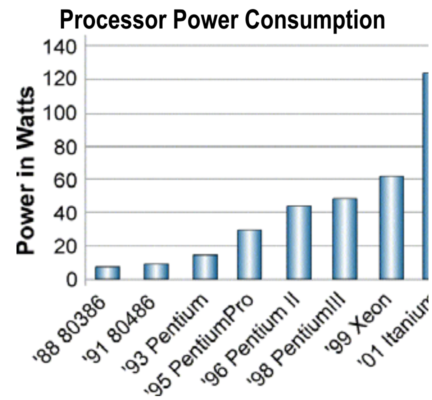

As processors operate faster and add more features, they consume more power and generate more heat. Processors have a safe temperature range for normal operation. If a processor operates beyond its safe temperature range, problems can vary from system crashes and random reboots to memory and application errors.

energy efficient power supplies save money

More efficient power supplies can lead to significant savings in power and cooling costs. As described previously, the Power Factor reflects how efficiently power supplies use electricity. PFC power supplies for servers usually have a power factor between 0.95 and 1.0.

active inrush current protection is required

Electronic devices containing solid-state power supplies experience an input current during initial start-up that can be several times greater than their operating current. High current surge during start-up, referred to as inrush current, can affect electrical systems by tripping fuses and circuit breakers unnecessarily. Given the number of servers in a datacenter, power supplies must contain inrush current surge protection to limit the current drawn during startup. If the power supply does not use current surge protection, relays and fuses must be rated higher than any possible surge current.

hot-plug capability for fault-tolerant systems

HP ProLiant servers in fault-tolerant service have two hot-plug power supplies. The power supplies share the load equally. If one power supply fails, the other supply immediately takes over the full load. Customers can replace the failed power supply while the other power supply is in use. When the failed power supply is replaced, the new power supply begins sharing the load shortly after being plugged in.

power and cooling planning

Given the trends in power consumption and heat density, plans for datacenter build outs and upgrades should include a modular design that provides sufficient headroom for increasing power and cooling needs. A modular design will provide the flexibility to scale capacity in the future when planned and unplanned changes become necessary. The initial cost of adding headroom to power and cooling design criteria may be insignificant compared to the cost of retrofitting the datacenter later. The budget for the project will determine the amount of design headroom that is possible. The following sections describe more specific considerations and requirements for power and cooling resources.

power consumption and heat load

Larger systems are moving to higher amperage or three-phase power. Many enterprise-class machines presently use three-phase power and most datacenters are already wired for three-phase power. A common and economical method of supplying power to high-density datacenters is to use a 208-V, three-phase system, known as high line AC power. The system operates with 208 V between any two transformer windings, giving 120 V to neutral (Figure 6). Existing datacenters with 110-V power will need to upgrade to high line power. In new datacenter build outs, HP recommends providing two 30-amp 208/230 single-phase circuits per rack. Future high-density server environments may require up to two 50-amp circuits per rack.

figure 6. High line AC power

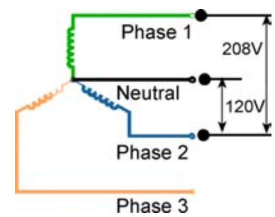

Although racks of equipment currently use single-phase power, it may be necessary to switch to three-phase power or higher amperage as technology changes. Customers should consider three-phase power when:

- the load exceeds 6 kVA
- a centrally-managed UPS solution is needed
- a single UPS solution to support the entire datacenter is warranted

Rack manufacturers provide wattage and BTU information in their product specifications. For precise facilities planning, HP provides a ProLiant Class, Rack/Site Installation Preparation Utility, accessible from the Active Answers web page. This utility includes a Power Calculator for each ProLiant server (Figure 7). ProLiant Power Calculators are macro-driven Microsoft Excel spreadsheets developed for two purposes:

- to review server loading to determine the number of power supplies required to provide redundant power supplies, and
- to approximate the electrical and heat load per server for facilities planning.

figure 7. ProLiant Class, Rack/Site Installation Preparation Utility available on the HP website

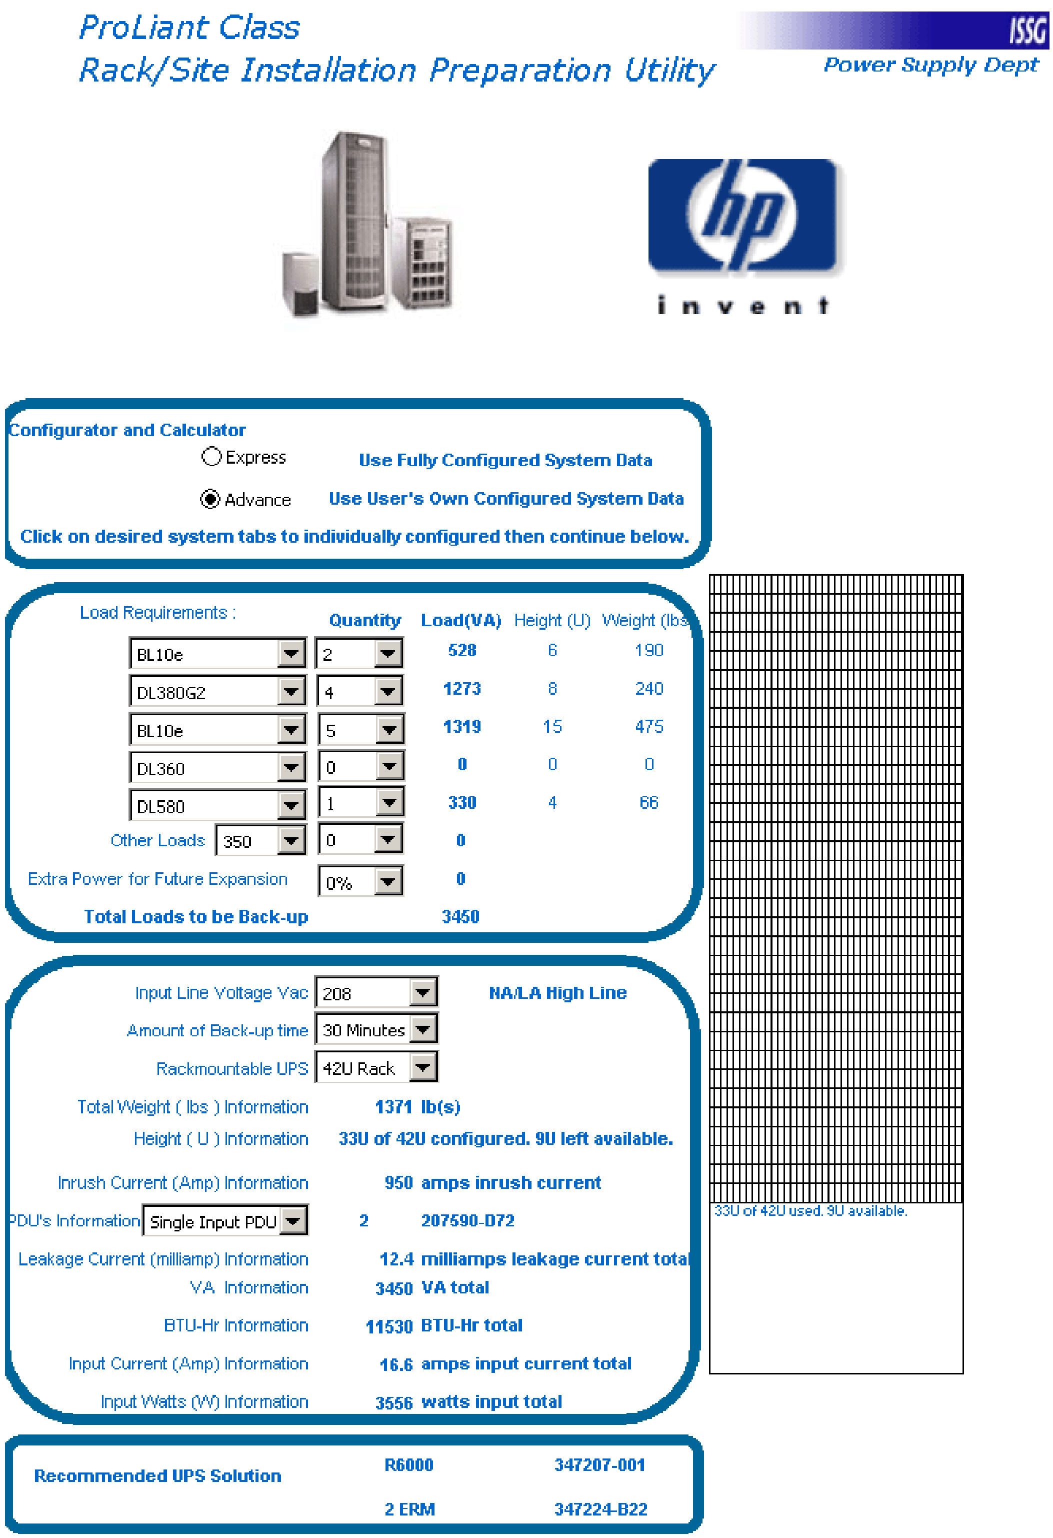

cooling requirements in high-density datacenters

This section summarizes several factors that should be considered when planning the cooling of datacenters with high-density heat loads. Whether designing a new datacenter or retrofitting an existing one, customers should work with knowledgeable heating, ventilating, and air conditioning (HVAC) engineers to ensure adequate cooling.

capacity of HVAC units

The heat load of the equipment determines the number and capacity of HVAC units required. The heat load of equipment is normally specified in BTU/hr. However, HVAC unit capacity is often expressed in tons of refrigeration, where one ton corresponds to a heat absorption rate of 12,000 BTU/hr. The "tons" capacity rating is measured at 80°F; however, the recommended operating conditions for HVAC units are 70° to 72°F and 50 percent relative humidity (RH). At 72°F, the HVAC unit output capacity is considerably reduced. Furthermore, the tons rating is very subjective because it is based on total cooling, which is comprised of "sensible cooling" and "latent cooling."

Computer equipment produces sensible heat only; therefore, the sensible cooling capacity of an HVAC unit is the most useful value. For this reason, HVAC unit manufacturers typically provide cooling capacities as "total BTU/Hr" and "sensible BTU/Hr" at various temperatures and RH values. Customers should review the manufacturer's specifications and then divide the sensible cooling capacity (at the desired operating temperature and humidity) by 12,000 BTU/Hr per ton to calculate the useable capacity of a given HVAC unit, expressed in tons of cooling.

Cooling capacity is also expressed in cubic feet per minute (cfm). The volume of air (in cfm) required is related to the moisture content of the air and the temperature difference between the supply air and return air (ΔT):

$$\text{Cubic feet per minute} = \text{BTU/hr} \div (1.08 \times \Delta T)$$

The capacity of each HVAC unit determines its effective cooling area. The effective cooling area takes into account the cooling footprint of the equipment.

cooling footprint

The floor area that each rack requires must include an unobstructed area to draw in and discharge air. Almost all HP equipment cools from front to rear so that it can be placed in racks positioned side-by-side and arranged in rows front-to-front and back-to-back to form alternating hot and cold aisles. The equipment draws in the cold supply air and exhausts warm air out the rear of the rack into hot aisles. The amount of space needed between rows of racks is determined by the cooling footprint of the equipment. The cooling footprint (Figure 8) includes width and depth of the rack plus the area in front for drawing in cool air and the area in back for exhausting hot air. Equipment that draws in air from the bottom or side or that exhausts air from the side or top will have a different cooling footprint. The total physical space required for the datacenter includes the cooling footprint of all the racks plus free space for aisles, ramps, and air distribution. Typically, a width of at least two floor tiles is needed in the cold aisles between racks, and a width of at least one unobstructed floor tile in the hot aisles is needed to facilitate cable routing.

Figure 8. Cooling footprint

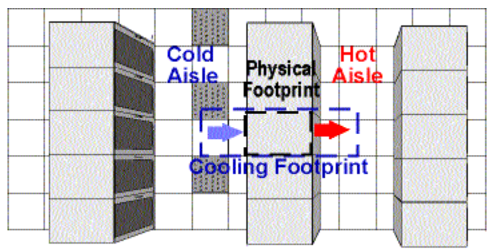

placement of HVAC units

The geometry of the room and the heat load distribution of the equipment determine the best placement of the HVAC units. HVAC units can be placed inside or outside the datacenter walls. Customers should consider placing liquid-cooled units outside the datacenter to avoid damage to electrical equipment that could be caused by coolant leaks. HVAC units should be placed perpendicular to the rows of equipment and aligned with the hot aisles (Figure 9). Rooms that are long and narrow may be cooled effectively by placing HVAC units around the perimeter. Large, square rooms may require HVAC units to be placed around the perimeter and through the center of the room.

Figure 9. HVAC units should be placed perpendicular to rows of equipment.

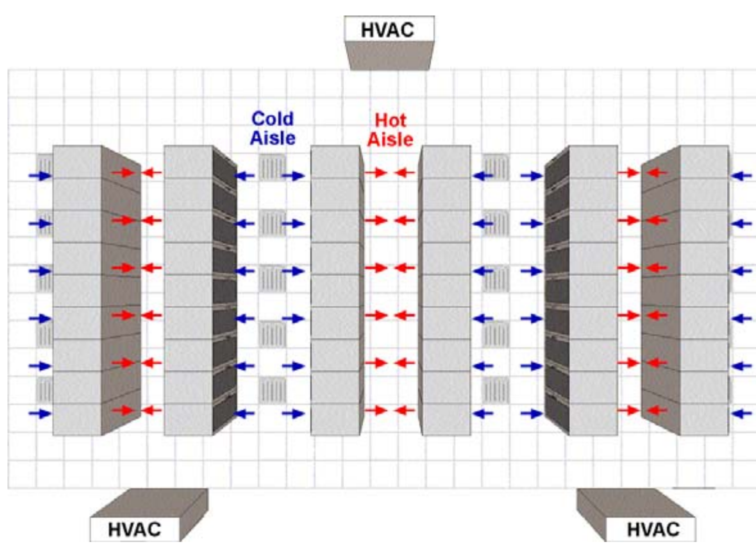

raised floors

High-density datacenters use a downward airflow pattern, which requires a raised floor configuration (Figure 10). Raised floors typically measure 18 inches (46 cm) to 24 inches (61 cm) from the building floor to the top of the tile. Floor tiles are usually 24 inches (61 cm) square and supported by a grounded grid structure. The main purpose of the plenum beneath the raised floor is to direct cool air to the racks. In a downward airflow pattern, air currents are cooled and heated in a continuous convection cycle. The HVAC unit draws in warm air from the top, cools the air, and discharges it into the supply plenum beneath the floor. The static pressure in the supply plenum pushes the air up through perforated floor tiles in cold aisles. Ideally, the warm exhaust air rises to the ceiling and flows along the ceiling back to the top of the HVAC unit to repeat the cycle.

figure 10. Airflow patterns for raised floor configurations

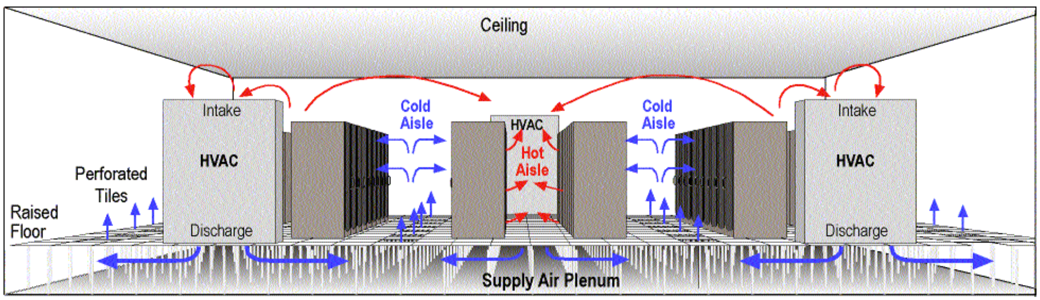

factors affecting airflow distribution

static pressure in the supply air plenum

For adequate airflow, the static pressure in the supply air plenum beneath the raised floor must be greater than the pressure above the raised floor. Typically, the plenum pressure should be at least 5 percent greater than the pressure above the floor.

The percentage and placement of perforated floor tiles are major factors in maintaining static pressure. Perforated tiles are classified by their airflow percentage. Airflow percentages vary from 25 percent (the most common) to 56 percent (for high airflow). A 25 percent perforated tile provides 548 cfm at a 5 percent static pressure drop, while a 56 percent perforated tile provides 2006 cfm. Tiles should be placed in front of at least every other rack. The high velocity discharge from the HVAC unit reduces the static pressure through perforated tiles nearest the unit, causing inadequate airflow (see Figure 11). The static pressure increases as the discharge moves away from the unit, thereby increasing the airflow through the perforated tiles. To counter this situation, airfoils under the raised floor can be used to divert air through the perforated tiles.[1] Another option is to use a fan-assisted perforated tile to increase the supply air circulation to a particular rack or hot spot. Fan-assisted tiles can provide 200 to 1500 cfm of supply air.

---

[1] From *Changing Cooling Requirements Leave Many Data Centers at Risk*. W. Pitt Turner IV, P.E. and Edward C. Koplin, P.E. ComputerSite Engineering, Inc.

figure 11. High supply air velocity reduces static pressure closest to the discharge.

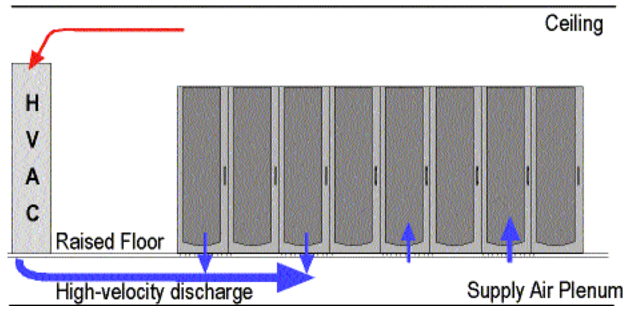

cooling range of HVAC units

The effective cooling range of each HVAC unit is determined by its capacity and the heat load of the equipment in its airflow pattern. The most effective cooling begins about 8 feet from the HVAC unit. Although units with capacities greater than 20 tons are available, the increased heat density of today's servers limits the cooling range to approximately 30 feet (Figure 12).

figure 12. Cooling ranges of HVAC units

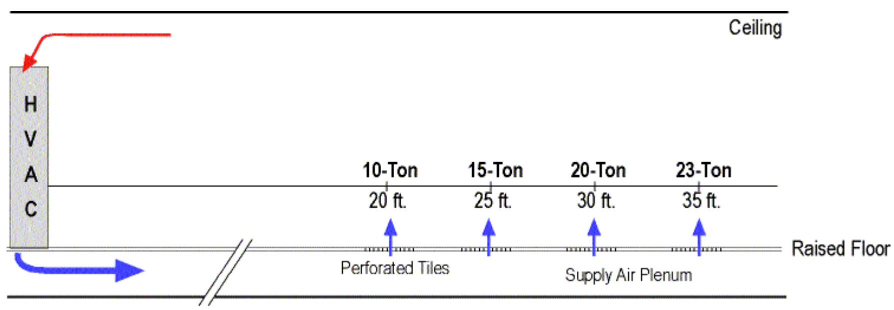

airflow blockages

The plenum is also used to route piping, conduit, and cables that bring power and network connections to the racks. In some datacenters, cables are simply laid on the floor in the plenum where they can become badly tangled (Figure 13). This can result in cable dams that block airflow or cause turbulence that minimizes airflow and creates hot spots above the floor.

U-shaped "basket" cable trays or cable hangers can be used to manage cable paths, prevent blockage of airflow, and provide a path for future cable additions. Another option is to use overhead cable baskets to route network and data cables so that only power cables remain in the floor plenum.

figure 13. Unorganized cables beneath raised floor

supply air plenum leaks

Electrical and network cables from devices in the racks pass through cutouts in the tile floor to wireways and cable trays beneath the floor. Oversized or unsealed cable cutouts allow supply air to escape from the plenum, thereby reducing the static pressure. Self-sealing cable cutouts are required to maintain the static pressure in the plenum (Figure 14).

figure 14. Self-sealing cable cutout in raised floor

airflow mixing

To achieve an optimum downward airflow pattern, warm exhaust air must be returned to the HVAC unit with minimal obstruction or redirection. Ideally, the warm exhaust air will rise to the ceiling and return to the HVAC unit's intake. In reality, only the warm air closest to the intake may be captured; the rest may mix with the supply air. Mixing occurs if exhaust air goes into the cold aisles, if cold air goes into the hot aisles, or if there is insufficient ceiling height to allow for separation of the cold and warm air zones (Figure 15). When warm exhaust air mixes with supply air, two things can happen:

- The temperature of the exhaust air decreases, thereby lowering the useable capacity of the HVAC unit.
- The temperature of the supply increases, which causes warmer air to be recirculated through computer equipment.

The following section describes some ventilation configurations that can help minimize airflow mixing.

figure 15. Mixing of supply air and exhaust air

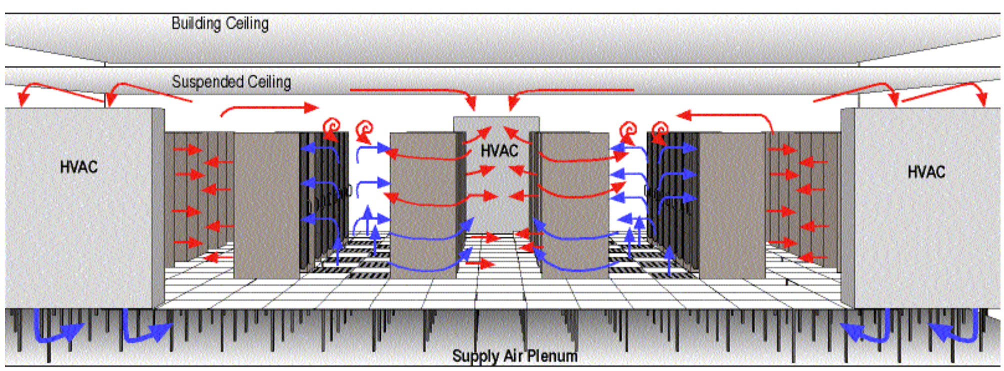

configurations for high-density datacenters

ceiling return air plenum

In recent years, raised floor computer rooms with very high heat density loads have begun to use a ceiling return air plenum to direct exhaust air back to the HVAC intake. As shown on the right of Figure 16, the ceiling return air plenum removes heat while abating the mixing of cold air and exhaust air. Once the heated air is in the return air plenum, it can travel to the nearest HVAC unit intake. The return air grilles in the ceiling can be relocated to best suit the layout of computer equipment.

figure 16. Ceiling return air plenum

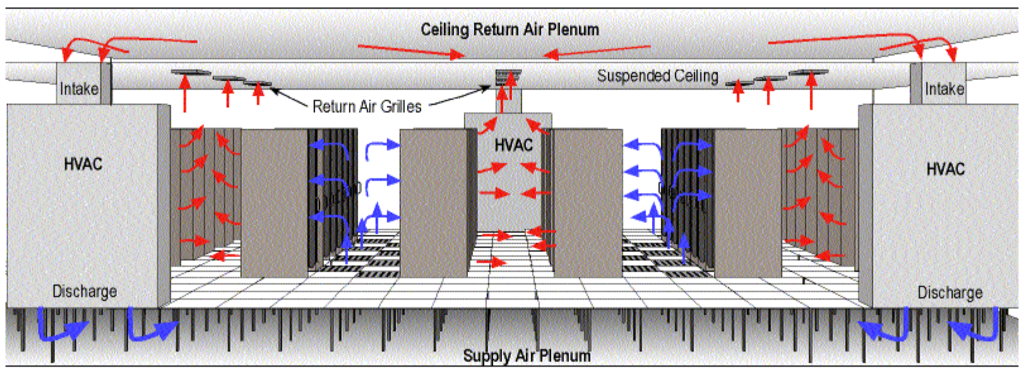

dual supply air plenums

As power and heat densities climb, a single supply air plenum under the raised floor may be insufficient to remove the heat that will be generated. High-density solutions may require dual supply air plenums, one above and one below (see Figure 17). In this configuration, additional supply air is forced downward in the cold aisle.

figure 17. Dual air supply plenum configuration for high-density solutions

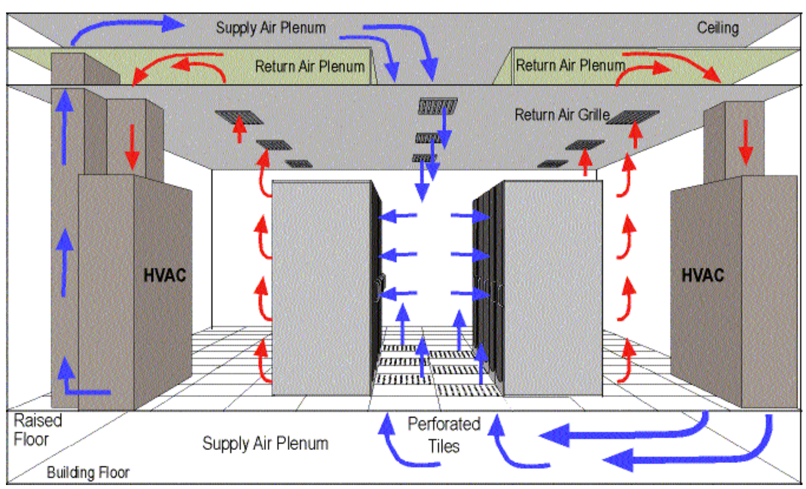

separate workspace and computing space

The high-density cooling solutions described above work well for equipment but do not work well for people. Customers should consider moving workstations outside of the datacenter and operating the datacenter as a "lights out" facility. This will permit lower operating temperatures and higher ambient noise.

conclusion

This paper described some of the challenges that must be addressed during the design of a new or retrofitted datacenter. The trends in power consumption of computing components necessitate modular datacenter designs with sufficient headroom to handle increasing power and cooling requirements. To determine these requirements, several factors should be considered, including the capacity and placement of the HVAC units and the geometry of the room. In addition, high-density datacenters require special attention to factors that affect airflow distribution, such as supply plenum static pressure, airflow blockages beneath raised floors, and configurations that result in airflow mixing in the datacenter.

feedback

Please direct comments regarding this communication to the ISS Technology Communications Group at this Internet address: TechCom@HP.com

10/2002

Document Number TC021002TB