

NEW DEVELOPMENTS IN data center design



CONNIE ZHOU / COURTESY GOOGLE

Google's data centers include this massive server room in Council Bluffs, Iowa. The company operates some 900,000 servers in about a dozen locations worldwide.

BY C.C. SULLIVAN AND BARBARA HORWITZ-BENNETT,
CONTRIBUTING EDITORS

LEARNING OBJECTIVES

After reading this article, you should be able to:

- + **DISCUSS** two primary strategies for free cooling for energy savings and environmental improvement in data centers.
- + **DESCRIBE** potential uses of renewable energy for data centers.
- + **EXPLAIN** the pros and cons of overhead vs. underfloor cabling and air supply in data centers, especially the implications for occupant safety.
- + **UNDERSTAND** trends in data use, computing, and facility infrastructure that are changing data center design, particularly packaged systems and cloud computing.

From the dozen or so facilities housing Google's 900,000 servers to the sprawling server farms of Facebook to Amazon's seven sites scattered around the world, today's data centers must accommodate massive power demand, high heat loads, strict maintenance protocols, and super-tight security. Among owners' top concerns are constant increases in storage and processing requirements, affecting rack densities, backup power needs, and MEP infrastructure.

The sheer numbers behind these data facilities are staggering. According to *WholsHostingThis.com* and the blog *Storage Servers*, Google's data centers alone use 260 million watts of power a year—enough to power 200,000 households. Facebook serves a user base that adds about seven petabytes—seven thousand trillion bytes—of photos to the platform every month.

Among the biggest players has been Microsoft, which has invested some \$23 billion on data centers and their contents to date. Last year, the company spent \$112 million on a single facility.

Building Teams can have a tough time keeping up with rapid change in this sector. What works today may quickly be obsolete, especially if the end-user's power requirements are doubling every nine months. However, oversizing is not necessarily the most viable option.

"In fear of what the future may bring, clients are now asking us to install as much MEP infrastructure capacity as is physically and financially possible," says John Yoon, PE, LEED AP ID+C, Senior Electrical Engineer with McGuire Engineers (www.mcguireng.com). "But blindly increasing MEP infrastructure capacity in an attempt to stay ahead of future IT equipment needs is not a financially sustainable solution in the long run."

Instead, the generally recommended strategy is more effective management and utilization. This involves finding a sweet spot that meets current and predicted capacity needs without grossly oversizing the mechanical plant, says Daniel Bodenski, Director of Mission Critical Services at AEC and consulting firm CRB (www.crbusa.com).

Data center experts classify **scalable designs and builds** as an important long-term solution. This concept involves master planning a full build-out, but constructing the first data module at a manageable size for the client's current needs.

"Design the infrastructure to be easily repeatable and rapidly deployable so that future modules can be added in a seamless fashion," says Charles B. Kensky, PE, LEED AP, CEA, Executive VP with Bala Consulting Engineers (www.bala.com). "This allows the end-user to add power and cooling capacity without disrupting ongoing operations."

The need for **uninterruptible service** is another critical planning consideration. Some clients try to manage risk with geographical diversity or remoteness; others prefer the opposite approach. "The financial sector has been locating their data centers for years in outlying areas rather than in the major metropolitan areas," says Bill Brody, VP with construction management firm B.R. Fries (www.brffries.com). "These are generally less high-risk areas, out in the countryside. Yet others have data centers right in the city, like tech centers and other companies that really want to closely monitor or centralize their operations."

Brody, who works with digital media companies and institutional R&D entities, says many tech companies and universities just lease their data centers from companies that own and operate facilities for multiple clients. "But Wall Street banks, for example, really want their data centers operated by their own people, so they can keep track of them and manage the security issues better," he notes.

MODULARITY, FREE COOLING, AND THE QUEST FOR LOW PUE

Once a client has selected the preferred location, the Building Team must devise effective designs. Modularity is a guiding principle. By plugging chillers, generators, and uninterruptible power systems (UPS) into a standardized supporting framework, clients can defer



COURTESY BALA CONSULTING ENGINEERS

Unlike air-conditioning systems placed around the perimeter of computer rooms, closed-couple cooling and containment systems are sited right next to the IT load. The cool air has a much shorter path to travel and can cool the cabinets more efficiently.

capital expenses to a just-in-time scenario. Rob Sty, PE, SCPM, LEED AP BD+C, a Principal in the Phoenix Technologies Studio of SmithGroupJJR (www.smithgroupjir.com), notes that this method also reduces **stranded capacity**: capacity that can't be used by IT loads due to system design or configuration problems.

Fortunately, the industry is seeing a higher level of coordination between data center managers and IT directors, which is enabling power and cooling systems to be more closely coupled with actual and predicted demand. Aaron Duda, PE, LEED BD+C, an Uptime Accredited Tier Designer (UATD) and Senior Associate with Environmental Systems Design (www.esdglobal.com), says that this partnership is allowing performance-based deployments instead of initial overdesign.

Beyond modular, scalable designs, more efficient power and cooling are a high priority. Teams are striving to leverage climate patterns, existing infrastructure, or both to gain **free cooling**—a strategy required for data centers by the latest energy codes. Free cooling uses outside air or an existing source of cold water to help chill buildings and equipment.

Nearly every data center designed by MEP consultant Syska Hennesy Group (www.syska.com) in recent years has included some form of free cooling, the most popular of which is the water-side economizer, according to James Coe, PE, RCDD, Director of Critical Facilities for the firm's Atlanta office. "If the end-user has a water-cooled chiller plant, this is the lowest-cost solution for free cooling, and it usually pays back quickly," he says.

Direct air-side economizers can be more expensive and may require a larger footprint; their filtration requirements and costs can run high, and they may not easily isolate ambient contaminants or handle weather anomalies, says ESD's Duda. The HVAC controls may be hard-pressed to achieve stable humidity conditions in a direct air-side system, he adds.



Gaining traction in the mission-critical marketplace: indirect outdoor-air economizers with evaporative cooling, which take advantage of ambient heat rejection while mitigating many of the problems associated with processing outdoor air, says Duda.

Sty reports that the University of Utah's data center is reducing central plant operating hours via air-side economizer operations, and is poised to achieve a **power usage effectiveness or PUE**—the ratio of the total energy used by a data center to the energy delivered to the computing equipment—of 1.25 at full load. The strategy will also save about 10 million gallons of water per year.

Yoon says Google's data centers have managed to push their PUE down to 1.2, an improvement over the industry average of 1.8 for large data centers.

The practicality of free cooling, regardless of method, is tightly linked to a facility's parameters for indoor temperatures—especially, appropriate temperatures for air entering servers via ventilation fans. ASHRAE Technical Committee 9.9—Mission Critical Facilities, Technology Spaces and Electronic Equipment has issued recommendations and thermal guidelines encouraging teams to consider **higher server inlet temperatures**, says Sty. Allowing higher temperatures can significantly reduce costs by increasing the number of hours possible for economizer operations.

“The basic premise is that server equipment is generally much more robust than we give it credit for, and it isn't necessary to cool a server room like a meat locker,” Yoon says. In 2008, Christian Belady, PE, Principal Power and Cooling Architect for Microsoft, performed a now-famous experiment wherein a rack of servers, protected from the elements by only a flimsy metal-framed tent, ran reliably for half a year.

Sometimes it's hard for data center managers to overcome the perceived risk of such a design approach. “While increasing the cold-aisle temperature set points may dramatically increase the feasibility

of air-side economizers in many regions, you have to convince the IT manager to embrace a design that goes against the traditional convention that low temperature and separation from the surrounding environment equals reliability,” says Yoon.

To help data center designers implement economizer systems and take advantage of free cooling, Bodenski recommends the following plan of action:

- Determine the potential number of free cooling days for both air-side and water-side economizer systems.
- Get aligned with the end-user's IT equipment requirements—supply air temperature and chilled water supply temperature—and deployment strategy. Will the IT load be deployed 100% on day one, or just 15% on day one?
- Determine the type of distribution system—underfloor, overhead, or other.
- Develop analyses for monthly electrical consumption and water consumption, as well as a payback comparison for air-side and water-side economizer options using net present value.
- Work with the end-users to obtain criteria to develop the engineering and economic analyses, notably geographic location, system type, IT load deployment, water-side economizer pumping penalty, air-side economizer fan penalty, utility rates, demand rate, water utility rate, maintenance costs, internal rate of return, utility and maintenance escalation, and construction cost inflation.
- Consider partial chiller loads, a hybrid air-side and water-side economizer, or both.
- Take into account the cost and complexity of modifying the base building architectural system and structural system to accommodate an air-side economizer.

ACTIVE COOLING STRATEGIES FOR ROWS, RACKS, AND CABINETS

Free cooling will take a data center only so far. Engineers still must deal with increasing rack densities and the need to prevent hot spots. This often means bringing the cooling to the load, not just the room, says Bala's Kensky. “Plan it and model it, but do it efficiently,” he says. “Higher density does not mean just adding cooling capacity. It requires an intelligent, efficient approach to delivering the cooling at the point of greatest need, and removing the heat effectively.”

Raised-floor systems work well for densities below 10 kW/cabinet. High densities require other solutions. For instance, central station air-handling units (AHUs), in lieu of computer-room air handlers, can employ a “ballroom” delivery of air into the cold aisle and integrated rack, says Sty. In-row cooling is also growing in popularity.

“As rack densities increase beyond the limits of what traditional air-cooled cabinets can support, water cooling at the cabinet may become more viable,” Sty suggests. He points to the High Performance Computing Data Center in the new Energy Systems Integration Facility at the National Renewable Energy Laboratory, in Golden, Colo., where cooling liquid is delivered directly to the cabinet. This strategy yields extremely high thermal efficiencies and contributes to a PUE of 1.06. Rear-door coolers—both water- and refrigerant-based—can

software company upgrades to save \$1 MILLION A YEAR ON ELECTRICITY

Recently, Bala Consulting Engineers helped a global business software company upgrade the infrastructure at a legacy 40,000-sf data center, achieving an annualized power usage efficiency of 1.43.

The upgraded MEP infrastructure provides hot-air containment via vertical exhaust ducts and solid rear-door cabinets for the existing cabinets. The system also incorporates an oversized cooling tower, high-efficiency chillers, and a heat exchanger for free cooling.

“The chilled water supply temperature was raised from 45 to 61°F, while maintaining a 75°F cabinet inlet temp, providing a warm water distribution system and increasing the number of free cooling hours per year,” says Charles B. Kensky, Executive VP with Bala.

UPS modules were swapped for high-efficiency units with a higher voltage distribution and high-efficiency transformers.

Though the data center increased its IT load from 1,200 kW to 4,000 kW, the lower PUE translates to an operating cost of \$0.09/kWh. The upgrades are saving the company nearly 13,000 MWh, or \$1.16 million in operational costs per year.

accommodate loads up to 30 kW/rack and are highly adaptable to retrofits, says CRB's Bodenski.

Yoon says top-of-row bus duct and in-rack liquid cooling can be great choices for high-density server lineups, but he also says they shouldn't be viewed as a one-size-fits-all solution. "Just because you can put 800-ampere bus duct above a 10-cabinet server lineup and also provide the cooling to make it work, should you?" he asks.

Yoon recalls one project for which an extensive blade server deployment was planned. A high-density, pumped refrigerant cooling system was installed to accommodate the load, only to have the client change its preferred server vendor. The blades were never fully deployed, and the load thresholds for effective operation of the cooling system were never reached. The cooling system had to be decommissioned and replaced less than five years later, according to Yoon.

To avoid such situations, McGuire Engineers generally recommends modular designs, specifying generic building blocks with a predefined per-cabinet power and cooling budget. "This entails considerable coordination with end-users so they understand the power and cooling budgets they have to live within," says Yoon. "Often, it's simply a case of redistributing IT equipment within the facility to even out temperatures and minimize the chance of hot spots."

New technologies are allowing more efficient data center layouts, such as 40- and 100-gigabit **top-of-rack (ToR) network architecture**, says Yoon. "This design methodology is much easier to implement, as ToRs have reduced the cabling requirements for high-bandwidth equipment," he says.

Some enterprise facilities are starting to consider **full-submersion cooling solutions**, which immerse the server electronics in a dielectric bath, to take densification to the next level. "The dielectric fluid is circulated over the server components, and then heat is rejected to an open-loop condenser water system," says ESD's Duda. "The dielectric can move more heat per volume unit, allowing the deployments to have a much smaller footprint."

HOT AISLES, COLD AISLES

Engineers are also striving to devise the best **aisle containment strategies**, preventing hot air exhausting from server cabinets from mixing with cold air supplied by the cooling system. Bala's Kensky recommends starting with simple strategies such as *brush grommets* (to seal cable openings in the floor) and *scalable blanking panels*, which fill unoccupied rack space to control airflow and enable servers to operate at a cooler temperature. End caps, aisle covers, and chimneys—vertical exhaust ducts extending from the top of cabinets to the ceiling plenum—can all be effectively deployed to contain hot air.

Cold-aisle-driven plans, which concentrate on maintaining temperatures by keeping cold aisles cold through various barrier strategies, work well for retrofits with in-row cooling or a raised-floor environment, saving on ductwork and construction expenses, says Bodenski. However, hot-aisle plans, which emphasize environmental control by confining and exhausting hot air, remain the preferred method for many new data center build-outs. Among the advantages:

- IT personnel can work in a 75°F environment, compared with 100°F or more in cold-aisle containment settings.
- End-users can take advantage of extended air-side or water-side economizer hours, which increases mechanical system efficiency and lowers PUE.
- Some IT equipment may not be able to survive outside of the contained areas and would need evaluation in a cold-aisle setting.

One approach Syska Hennessy likes to pursue is **partial containment**, where the air pressure in the hot aisle and cold aisle are kept at the same levels. "With full containment deployed, the pressure in the hot aisle can be higher than the cold aisle," Coe explains. "The server fans may have to be sped up and draw more critical power to overcome this pressure."

Aisle containment is not a simple spec. The concept has evolved from a straightforward return air pathway to a sophisticated architectural solution supporting numerous infrastructure components, making full Building Team collaboration increasingly important. "Coordination

PLUG AND PLAY: planning for modular expansion

With data centers growing at such a rapid clip, coupled with the reality of high O&M costs, most end-users can only afford to support current loads—even when they know loads will most likely increase very soon. Plug-and-play infrastructure is a response to the problem: complete systems are shipped in containers, ready to go. Lead times are quick, and testing and commissioning can be conducted in an isolated fashion. Plugging a new module in should have little impact on critical power and cooling—as long as modularity was planned at the outset, with appropriate electrical and mechanical distribution systems available.

Plug-and-play infrastructure works particularly well for data centers that lease space to contracted

clients, says Rob Sty, Principal in the Phoenix Technologies Studio at SmithGroupJJR. "It is a very economical approach to deferring capital expense costs until a specific customer has been identified, and the best example of scalability," he says. "The benefit is that a cost estimate on the next phase of infrastructure can be easily defined, and the timeline of construction can be significantly reduced." The difficult part comes in correctly predicting the right size of the module, and client requests for customized solutions that do not fit into that model.

Authorities having jurisdiction may be skeptical of plug-and-play additions. "While the AHJ may not have an issue with a single piece of equipment, cramming computer-room air-conditioners, UPS, fire-suppression systems, and equipment racks

in an intermodal container will always cause an increased level of scrutiny," says John Yoon, Senior Electrical Engineer at McGuire Engineers.

One other downside of this deployment strategy is obsolescence, according to Aaron Duda, Senior Associate at Environmental Systems Design. "The information technology sector is quick to change and adopt new processing strategies and hardware components to give users a seamless experience, while plug-and-play infrastructure relies on proven and dependable technologies that provide a stable and unchanging platform," he says. "As IT deployments change and allow for more aggressive environmental designs, a plug-and-play deployment may not allow the facility to take full advantage of new options."

with all design disciplines is necessary to ensure that the pod containment architecture fits into the entire building system,” says Sty.

To make it easy to set up containment solutions, some vendors are now offering prefabricated systems as a turnkey design, manufacturing, installation, and commissioning package.

GREENER POWER GENERATES DEBATE

Until recently, data centers’ enormous power and water use has escaped widespread criticism; generally, the consumption has been viewed as a necessary evil in a booming industry. But such governmental and public tolerance is likely to change. “It is only a matter of time before the environmental impact of these types of facilities is put on display,” predicts ESD’s Duda. “Many enterprise users fear this bad publicity so much that they are investing in on-site photovoltaic or fuel cell installations to offset their grid use of power. They’re also examining on-site use of process water from the cooling systems, to curtail the discharge of the fluid to storm or sanitary systems.”

Though fuel cells aren’t cheap, they can make sense in areas where utility rates run high and federal and state tax incentives are available. Kensky reports that fuel cells by one manufacturer (Bloom Energy, Sunnyvale, Calif.) can run on a variety of inputs and withstand high temperatures. eBay is using these natural gas-powered cells to power their newest data center in Salt Lake City, even when regular utility services are available. They have opted not to install generators or UPS equipment. For some projects, the waste heat from fuel cells and microturbines can be captured in **combined-cooling-and-power systems**, driving absorption chillers to contribute to cooling.

Solar remains the low-hanging fruit among renewable power solutions for data centers. Even though facilities may require acres of PV panels to achieve a reasonable result, experts say the investment can make sense. Wind power also requires lots of space. “If I had to choose an alternative, I would want to go hybrid: a hydroelectric base supplemented with solar, wind, or both, with generator or utility backing,” says Chris McLean, PE, Director of Data Center Design at Markley Group (www.markleygroup.com).

Designers are also trying to ratchet down power use by delivering DC power to the cabinets and specifying higher distribution voltages. When DC power is supplied to IT equipment, AC power conversion (long a standard aspect of data center engineering) is no longer necessary; this yields a 2-3% efficiency gain and can obviate the need for a UPS inverter section and downstream power distribution units, says Yoon.

Higher delivery voltages— for example, 230/400VAC and 240/415V—are another hot topic in the data center design community. They’re still relatively rare in North America but can be considered in high-density installations, according to Yoon. This strategy enables equivalent power delivery at a relatively low current and minimal transformer losses. However, there is a higher risk of arc flash that could injure inadequately trained personnel, particularly with top-of-row busways. Yoon warns that the relative rarity of these alternate power-delivery methods also carries an inherent cost in the form of proprietary server equipment and specialized operational methods.

Even higher voltage—380V DC—is on the far horizon.

CABLING: OVERHEAD VS. UNDERFLOOR

Experts debate whether it’s better to place data and power cabling overhead or under the floor. Bala’s Kensky says overhead trays deliver a number of advantages. They make it easy to implement additions and changes, offering greater long-term flexibility. Putting cable high also minimizes future disruption to access-floor cooling plenums.

When combined with overhead ducted supply air to the cold aisle, *overhead cabling* can allow teams to forgo expensive raised-floor systems altogether. Cabinets can be placed directly on the structural slab. Power and low-voltage systems can coexist in a fully coordinated modular design.

Perhaps the biggest benefit, says Syska Hennesy’s Coe, is that if power and data cables are placed overhead, and there is no raised floor (or the raised floor is only used to supply cold-aisle air), an emergency power-off (EPO) push button is no longer required, per the National Electrical Code.

“These EPO buttons are typically located at every IT room exit,” Coe explains. “They de-energize IT equipment in the room and stop the supply of underfloor air. However, EPOs make IT managers very nervous since outages are bad for their job security.”

An often-overlooked benefit of overhead installation, says Coe, is better O&M practices. When cabling is hidden below the floor, some



COURTESY SMITHGROUPJJR

For greater flexibility, all power and data cabling can be supplied overhead, as in this data center installation at the University of Utah, designed by SmithGroupJJR. Take care not to block the path of lighting.

installers will just abandon decommissioned infrastructure, creating blockages over time. “When cabling is visible overhead, it will typically be installed with better workmanship and be removed when it is no longer needed,” he says.

Duda points out that overhead power and data do require close coordination of lighting and environmental conditions. “Unencumbered light should be provided where equipment requires regular access,” he says. “If not properly planned, overhead cable trays can reduce the amount of light getting to the service spaces.” Designers need to be familiar with the Telecommunications Industry Association Standard TIA-942, Telecommunications Infrastructure Standard for Data Centers. The standard stipulates an 18-inch clearance from the top of cable pathways to sprinklers, which can be difficult to achieve in a legacy facility, or even in a new building with height constraints.

Overhead installation has other disadvantages. Vertical supports and tray locations can limit cable installation space. Ladders are required to access the cable. Unprotected infrastructure can be more vulnerable to damage, compared with underfloor installations. Cable ampacity (current-carrying capacity) may be compromised if installed over hot aisles with elevated temperatures, according to CRB’s Bodenski.

Underfloor layouts offer easy access for installation and removal by means of a simple floor tile-pulling device. “Cable deployments are hidden from view and protected by the raised-floor system, providing a clean aesthetic,” says Duda. “Properly planned underfloor deployments will have cavities large enough to accommodate cabling and provide clear pathways for air delivery.” Also a must: proper cable-management strategies that support airflow.

Underfloor layouts enable designers to reduce “white space” above the computer equipment and can make system coordination simpler. But they can also waste energy due to cables blocking the delivery air path, supply-air leakage from cable cutouts, and bypass-air leakage from power distribution unit cutouts, says Bodenski. UFADs also require a greater floor depth to accommodate cable trays, and trays installed in the supply airflow plenum will complicate maintenance.

Because each project is different, SmithGroupJJR’s Sty recommends weighing variables such as room height, rack and cabinet densities, and the overall facility cooling strategy before selecting a cabling location for any data center.

PACKAGED IT SYSTEMS OFFER INTEGRATED FACILITY SOLUTIONS

Following the trend toward greater systems integration, some data centers are moving to **converged infrastructure** to capture efficiencies and savings. Defined as multiple IT components—servers, data storage devices, and networking equipment—consolidated into a single package, convergence offers a reduced cabinet footprint, lowered power consumption, and an extended data center life cycle.

“Bringing together server storage, networking, and virtualization into an integrated solution that is managed as a single entity optimizes the IT infrastructure,” says Sty. “This has a direct impact on the facility’s supporting mechanical and electrical systems. Operational expenses can be reduced significantly through this operating platform.”



COURTESY SMITHGROUPJJR

In this example of hot-aisle containment at the University of Utah, central station air-handling units supply cold air directly into the cold aisle, while barriers keep the hot air from mixing with the cold. The design eliminates the need for raised-floor air distribution.

A recent market research study by International Data Corp. (www.idc.com) projects that overall spending on data center converged systems will grow at an annual rate of nearly 55%, reaching \$17.8 billion by 2016. By then it will account for 12.8% of the total storage, server, networking, and software market.

Duda says that clients with a greater appetite for risk are actively embracing convergence, while more conservative companies are likely to go with proven technologies. Coe adds, “These solutions are becoming more common, but rarely do we see an IT manager who will deploy a converged system alone. They generally deploy a variety of solutions.”

Yoon anticipates that the industry *will* move toward convergence. While data center managers, particularly in smaller facilities, traditionally focus on the IT equipment deployment and management, they will eventually be forced to pay more attention to MEP infrastructure, he says. Convergence can enable managers to take a systemwide approach.

A newer trend is **cloud computing**, which is transforming the process of computing from a product to a service. End-users are moving various aspects of operations—especially storage—from local servers to a network of remote servers hosted on the Internet.



mission-critical facilities

AIA CONTINUING EDUCATION

At \$200,000 a pop for a single on-site storage-area network cabinet, transferring data to a remote site can make lots of financial sense.

“Storage requirements are growing faster than almost any other sector of IT platforms,” says Kensky. “Even as storage technology becomes more efficient, we live in a smartphone and tablet world, and the demand for the cloud and storage will only grow.”

In healthcare, for example, new processing-intensive clinical technologies like computer-assisted diagnostics and telemedicine are combined with an ongoing need for secure medical-record storage. The scenario requires increasingly robust storage and processing protocols. Clients see the cloud as a way to help handle the enormous load, says Sty.

AEC firms like Syska Hennessy have started using the cloud themselves for archiving, e-mail, and some other IT applications. “The upside for firms of all sizes is that they can reduce their owned and maintained IT equipment and the need to provide connectivity, power, and cooling for less-critical IT functions,” says Syska’s Coe.

Cloud computing also allows data center operators to reduce the needed level of in-house infrastructure reliability. Cloud providers can shift IT activity to a different data center—or a different area within the same data center—in the event of an outage or failure. Designing data centers to a lower Tier level (a measure of uptime and redundancy) reduces capital and operating costs, since lower levels mean fewer generators, UPS modules, and chillers.

Cloud computing further enables operators to better standardize equipment and critical infrastructure across their portfolios, lowering costs and driving efficiency, says Bodenski.

Cloud computing fills an important niche, but it also raises issues of net neutrality, cybersecurity, and information ownership. Where confidentiality is a high priority, such as for critical operations and transactions, firms will most likely continue to want their own processing and storage infrastructure, says Duda.

MORE EFFICIENT PROJECT DELIVERY

Cloud or no cloud, technological evolution will continue to require smart thinking and nimble action from Building Teams and their clients. “The demand for data center critical load will grow until we have fully tapped the potential of computers, smartphones, tablets, Internet televisions, and whatever comes next,” says Coe. “That could be a while.”

Coe anticipates that the data center sector will eventually demand that Building Teams offer a more efficient construction-delivery process, pressed by the need to increase capacity, improve sustainability, and hold the line on cap-ex and operating costs. +

> EDITOR’S NOTE

Additional reading required for this course!

To earn **1.0 AIA CES HSW learning units**, study the article carefully and take the exam posted at www.BDCnetwork.com/datacentertrends



data centers AIA/CES MODULE

To earn **1.0 AIA CES HSW Discovery learning units**, go to www.BDCnetwork.com/datacentertrends and take the following exam.

- Energy codes now require that data centers incorporate free cooling, requiring:
 - Use of water-side economizers
 - Use of air-side economizers
 - Use of underfloor air delivery systems
 - Use of outside air or an existing source of cold water to help chill buildings and equipment
- Google’s data centers have a low power usage efficiency (PUE) of about 1.2, versus an industry average PUE of about
 - 1.8
 - 2.8
 - 18.0
 - None of the above
- True or false: Full-submersion IT deployment cooling solutions locate server systems in a geothermal well to reduce overheating.
 - True
 - False
- Which of the following cooling approaches is NOT a practical, viable option for high-density rack configurations in today’s data centers?
 - “Ballroom” delivery of air into cold aisles
 - In-row cooling
 - Central station air-handling units
 - Raised-floor systems
- Per the National Electrical Code, if both power and data connectivity are provided overhead and the raised floor is either eliminated or only used for supply cold-aisle air, then the emergency power-off (EPO) pushbutton is:
 - Recommended
 - Not required
 - Mandatory
 - Added automatically to the system
- The Telecommunications Industry Association (TIA) standard TIA-942, Telecommunications Infrastructure Standard for Data Centers, stipulates a clearance requirement from the top of the cable pathways to sprinklers of:
 - 6 inches
 - 12 inches
 - 18 inches
 - At least 4 feet
- Cloud computing allows data center operators to reduce the needed level of infrastructure reliability, because cloud providers can shift IT activity to a different data center during an outage or failure. Building to a lower Tier level reduces construction and operating costs because:
 - Fewer generators, UPS modules, and chillers are needed
 - Sprinklers and fire suppression are not required
 - They do not require cooling
 - None of the above
- Raised-floor infrastructure has many benefits for data center design, such as reduced ceiling heights, but it also raises the issue of:
 - Increased needs for BIM
 - Potential blocking of air delivery due to cables
 - Use of underfloor fire-suppression systems
 - None of the above
- True or false: Some data centers use electrical systems that are 230/400VAC and 240/415V, and even higher voltage—380V DC—is being developed for future use.
 - True
 - False
- Plug-and-play systems for data centers are available as containerized solutions, which provide for:
 - Shorter lead times
 - Simplified testing and commissioning
 - Minimal impact on critical power and cooling
 - All of the above

DATA CENTER BEST PRACTICES at the University of Utah

SmithGroupJJR was recently tasked with applying industry best practices to ramp up energy- and water-use efficiencies for 15 data centers and IT rooms operated by the University of Utah in Salt Lake City.

The local climate is cool and dry most of the year but includes extreme winter conditions and a short summer period with high wet-bulb temperatures. The design team created a psychrometric chart, based on weather data, mapping cooling distribution hours and recommended cooling strategies to maintain data floor setpoints.

The team determined that a predominately air-side economizer, full and partial approach, was the best option 73% of the year. A water-side economizer would work best 22% of the year, and vapor compression cooling would be required 5% of the year.

The final system design uses central station air-handling units with fan matrix air delivery directly into the room. Removal of raised-floor systems reduced system pressure drop and fan horsepower. Both the Enterprise (Tier III) and High Performance Computer (Tier I) areas use an air-side economizer with hot-aisle containment strategies, supporting IT loads of 20 kW per cabinet. Incorporating hot-aisle containment allowed higher cold-aisle temperatures, which increased the number of hours for partial or full economizer operation.

Air-side cooling is provided by multiple fan array units blowing directly into the cold aisles, which are kept at 76°(±2°)F. Fan speed is controlled by maintaining a slightly positive differential pressure set point between the hot and cold aisles. To block particulates that might be introduced through economizer operation, two layers of filtration were provided: MERV 8 prefilters and MERV 13 final filters. Outside air conditions are monitored by a weather station, guiding the choice of full or partial economizer operation.

During periods when the air-side economizer isn't a practical solution, the university's two-stage central cooling plant provides chilled water to the

indoor air-handling units. Due to the elevated temperatures supplied to the data floor, chilled water is delivered at 65°F with a 30°F temperature rise. Stage 1 of the central plant uses a water-side economizer via closed-circuit fluid coolers, which act as a dry cooler when dry-bulb temperatures are below 45°F. As the atmospheric wet-bulb temperature increases, or during an interruption of domestic water to the site, three air-cooled chillers operate in series to provide the required chilled water temperature. AHU coil selections were made to provide cooling at a 65°F chilled water supply temperature, increasing hours of water-side economizer use.

With power feeds from two substations, Phase 1 of the project installed three 2.5 MVA transformers, with room for two more in the future. Critical IT power was routed separately from HVAC equipment power and normal power, to assist with accurate PUE monitoring. HVAC equipment was sized in a modular fashion to support the transfer of load from one distribution path to another in the event of electrical failure. Enterprise-critical IT loads were separated from high-performance computer IT loads.



COURTESY SMITHGROUPJJR

Large air-handling units provide filtered, outside air for cooling this University of Utah data center 73% of the year, eliminating the need for chilled-water cooling. The latest energy codes require implementation of free cooling strategies in data centers.

Power from the centralized power distribution units is supplied to cabinets via overhead bus ducts, which support 120/208V or 480V distribution. The power strategy includes distributed redundant UPS with catcher architecture, featuring double-conversion high efficiency (96%) with maintenance bypass. This architecture allows for N+1 redundancy and improved efficiency through higher loading.

— ROBERT F. STY, PE, SCPM, LEED AP, SMITHGROUPJJR