High Efficiency Indirect Air Economizer-based Cooling for Data Centers

White Paper 136

Revision 1

by   Wendy Torell

Executive summary

Of the various economizer (free cooling) modes for data centers, using fresh air is often viewed as the most energy efficient approach. However, this paper shows how indirect air economizer-based cooling produces similar or better energy savings while eliminating risks posed when outside fresh air is allowed directly into the IT space.

Schneider Electric

Introduction

Aside from IT consolidation, the biggest opportunity for energy savings comes from the cooling plant, which, in many data centers consumes the same as, or even more than, the IT loads. The key to reducing cooling plant energy is to operate in economizer mode whenever possible. When the system is in economizer mode, high-energy-consuming mechanical cooling systems such as compressors and chillers can be turned off, and the outdoor air is used to cool the data center. There are two ways to use the outdoor air to cool the data center:

- Take outdoor air directly into the IT space, often referred to as "fresh air" economization
- Use the outdoor air to indirectly cool the IT space

The pro's and con's of each will be discussed later in the paper. There are several ways of implementing the second method, which are largely distinguished by how many heat exchanges occur between the indoor air and outdoor air. White Paper 132, *Economizer Modes of Data Center Cooling Systems,* compares economizer modes best suited for data centers.

This paper will illustrate why a cooling system with the following design principles reduces energy consumption by 50% and offers the flexibility and scalability needed for large data centers.

Design principle 1: Economizer mode is the primary mode of operation

Design principle 2: Indoor data center air is protected from outdoor pollutants and excessive humidity fluctuations

Design principle 3: Onsite construction time and programming is minimized

Design principle 4: Cooling capacity is scalable in live data center

Design principle 5: Maintenance does not interrupt IT operations

Figure 1 illustrates a cooling approach, referred to as indirect evaporative cooling, that, when packaged in a standardized self-contained footprint, meets these five design principles.

Figure 1

Cooling approach based on five key design principles
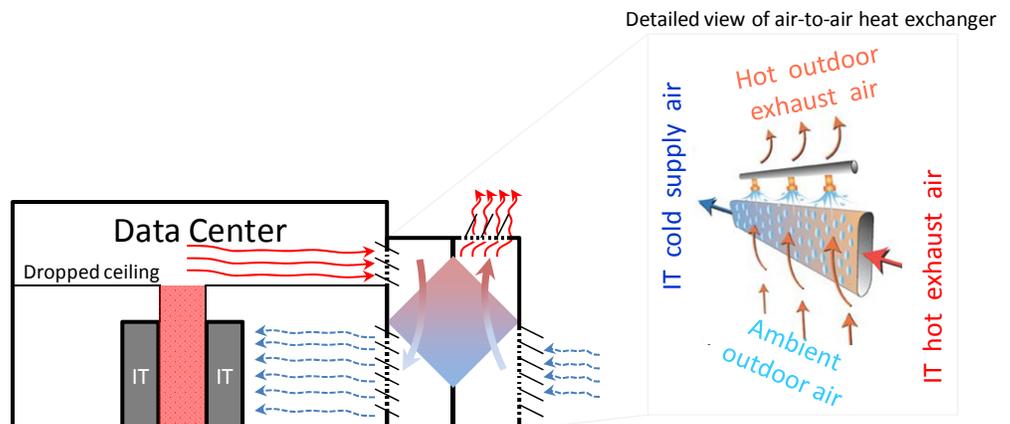
Economizer mode as primary mode of operation

The cooling approach of **Figure 1** maximizes economizer mode operation by reducing the number of heat exchanges to one with an air-to-air heat exchanger and by incorporating evaporative cooling. Alternatively, this design principle can be achieved with a fresh air (direct air) system which eliminates heat exchanges altogether.

Indoor data center air is protected from outdoor pollutants and excessive humidity fluctuations

Because the cooling method of **Figure 1** indirectly cools the air, the outdoor pollutants and rapid swings in temperature and humidity are isolated from the IT space. Alternatively, high quality filters can be implemented in direct (fresh) air systems to protect from outside contaminants and its control system can ensure the plant switches to backup cooling modes when dynamic weather changes occur beyond the data center's limits. Other indirect cooling architectures, as described in White Paper 132, *Economizer Modes of Data Center Cooling Systems,* can also achieve this design principle, but not typically while still maintaining economizer mode as its primary mode of operation (design principle 1).

Onsite construction time and programming is minimized

A cooling plant with integrated pre-programmed controls in a standardized self-contained system[1] allows for onsite construction and programming of the cooling plant to be reduced significantly. It also ensures reliable, repeatable, and efficient operation. As the data center industry continues to shift towards standardized modules (containers), this design principle will be achieved by many systems. White Paper 116, *Standardization and Modularity in Data Center Physical Infrastructure*, and White Paper 163, *Containerized Power and Cooling Modules for Data Centers*, discuss in greater detail how factory assembly and integration are driving down the amount of onsite construction and programming time.

Cooling capacity is scalable in live data center

With the dynamic loads that are characteristic of so many data centers today, it is critical that the cooling infrastructure can scale as the load scales. This can be achieved with "device modularity" as well as "subsystem modularity", as described in White Paper 160, *Specification of Modular Data Center Architecture*.

Maintenance does not interrupt IT operations

Redundancy (either through the use of internally redundant cooling modules within the system or the use of multiple systems) can eliminate single points of failure, and create a fault tolerant design enabling concurrent maintainability. In addition, a cooling system located outside or on the roof ensures maintenance activity takes place outside of the IT space which reduces the chance of human error impacting IT operations.

---

[1] A self-contained cooling plant is a complete cooling plant that is not dependent on other components to cool the data center.

Indirect vs. direct fresh air economization

As previously stated, design principle 2, *Indoor data center air is protected from outdoor pollutants and excessive humidity fluctuations,* can be achieved with both indirect and direct fresh air economizer-based systems. However, there are some key differences between the two approaches.

Using fresh air directly to cool a data center is often viewed as the most efficient economizer-based cooling approach. In general, lowering the number of heat exchanges is beneficial to the number of economizer mode hours and ultimately increased efficiency. And since direct air systems simply filter the outside air into the IT space, it has no heat exchanges (although the filters do result in added fan energy). For those data centers willing to let their IT environments experience a wide range of temperature and humidity conditions, this cooling approach is often the most efficient. However, today the majority of data center managers are risk-averse to higher temperatures and rapid changes in temperature and humidity. With rising densities and the adoption of containment practices, letting IT equipment run at higher temperatures increases anxieties because of what could happen if a cooling failure were to occur. When temperature & humidity thresholds are kept within ASHRAE recommended limits (discussed later), indirect air economizers actually provide greater efficiency than direct fresh air, in many geographies.

In addition, there continues to be reliability concerns over contaminants such as dust, chemicals from spills, smoke / ash, etc. Chemical sensors and filters can help protect against this but filters would need to be checked frequently, as clogged filters can prevent cooler air from entering the space, leading to thermal shutdown. Also, these filters result in an additional pressure drop on the air delivery system, which means more energy is required to move the same amount of air.

Over time, if the reliability and tolerance of IT equipment continues to improve, and if data center operators overcome the psychological barrier of requiring a tightly controlled environment, the use of direct fresh air systems may become more commonplace. White paper 132, Economizer Modes of Data Center Cooling Systems, provides further advantages and disadvantages of direct and indirect air economizer-based cooling.

An improved cooling approach

A system that addresses the five design principles described above is a self-contained cooling system with three modes of operation.

1. **Air-to-air economization** – takes the two air streams through an air-to-air heat exchanger – colder outdoor air, and the hotter indoor air that's heated from the IT equipment never mix.
2. **Air-to-air economization with evaporative cooling** – When the outdoor air isn't cool enough, evaporative cooling lowers the surface temperature of the heat exchanger through adiabatic cooling.
3. **Direct expansion (DX) cooling** – Worse case, when the air is too hot or too humid to support the IT inlet set point, DX cooling supplements either economizer mode.

The hot IT air is pulled into the module, and one of two modes of economizer operation is used to eject the heat. Based on the load, IT set point, and outdoor environmental conditions, the system automatically selects the most efficient mode of operation. The indirect air-to-air economization mode uses an air-to-air heat exchanger to transfer the heat energy from the hotter data center air to the colder outdoor air. When evaporative cooling is used, water is sprayed over the heat exchanger to reduce the surface temperature of the exchanger. This mode of operation allows the data center to continue to benefit from economizer mode operation, even when the air-to-air heat exchanger alone is unable to reject the data center

heat. The proportional DX mode provides additional cooling capacity when either economizer mode cannot maintain the IT inlet set point.

Figure 2 illustrates the application of this cooling approach in a new data center[2]. The cooling modules are placed outside of the facility – either mounted on concrete pads (as illustrated) or on the roof assuming it has appropriate weight bearing capacity.

Figure 2

Applications of cooling approach (perimeter of building and rooftop)
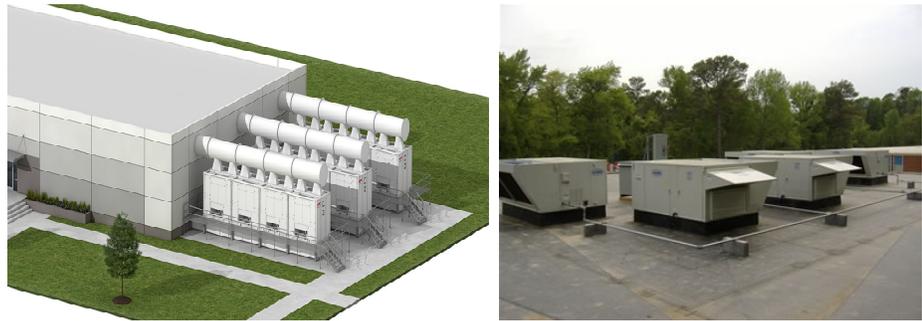
Efficiency improvement

With new guidelines and local regulations around the use of economizer modes, efficiency is a focal point for new data center design. To ensure the most efficient and effective form of cooling throughout the year, the cooling plant must use energy-saving economization as its primary mode of operation that maximizes localized climate conditions. In contrast, economizer mode operation in traditional designs have generally been perceived as an "add-on" to their primary cooling plant – to assist the high-energy-consuming mechanical plant when possible. Five factors, when combined, significantly improve the efficiency of a cooling system:

- Minimal number of heat exchanges between outdoor air and IT inlet air
- the use of evaporative cooling
- wider range of acceptable air temperature and humidity set points for IT equipment
- efficient components
- controls programmed at the factory

Impact of number of heat exchanges on economization
The more "heat exchange handoffs" that take place in economizer mode, the smaller the number of hours in economizer mode. **Figure 3** compares the heat exchange handoffs of a traditional chilled water cooling architecture with a plate-and-frame heat exchanger to a self-contained system with an air-to-air heat exchanger. Three heat exchanges take place in economizer mode for the traditional design – the cooling tower, the plate and frame heat exchanger, and the air handler, whereas the cooling design with an air-to-air heat exchanger has just the one exchange.

---

[2] The first example illustrates Schneider Electric's EcoBreeze modular indirect-air economizer cooling modules

Figure 3

Cooling architecture impacts economizer hours

Top – traditional chilled water plant (3 heat exchanges)

Bottom – self contained system (1 heat exchange)

Assumption: 100% load, St. Louis, MO, USA

> **> Mean Coincident Wet Bulb**
>
> Mean coincident wet bulb (MCWB) temperature is the average of the indicated wet bulb temperature occurring concurrently with the corresponding dry bulb (DB) temperature.

As this **Figure 3** illustrates, to obtain an IT inlet temperature of 70F (21.1C) on full economizer operation, the traditional design requires an outdoor air dry bulb temperature of 43.7F (6.5C) or lower and a mean coincident wet bulb temperature of 39.6F (4.2C) or lower. However, the air-to-air heat exchanger can achieve the same IT inlet temperature on full economizer operation with outdoor temperatures up to 59.3F (15.2C) and a mean coincident wet bulb of 53.5F (12C). This means there's an additional 16 degrees that the economizer mode can operate. For St. Louis, MO, those additional 16 degrees (from 43.7F to 59.3F) represent an additional 1,975 hours or 23% of the year.

Evaporative cooling

Evaporative cooling is another advantageous characteristic of high efficiency cooling modules because it increases the use of economization mode for many geographies, especially hot, dry climates. The energy benefit of evaporative cooling increases as the temperature difference between the ambient dry bulb and wet bulb temperatures gets larger.

Figure 4 illustrates how evaporative cooling can be implemented with an air-to-air heat exchanger. In evaporative cooling mode, water is sprayed evenly across the outside of the heat exchanger. As ambient air is blown across the outside of the heat exchanger (indicated by the number 1), the water evaporates causing a reduction in the outdoor air temperature[3]. The lower outdoor air temperature is now able to remove more heat energy from the hot data center air flowing through the inside of the tubes in the air-to-air heat exchanger (indicated by the number 2).

---

[3] Water requires heat to evaporate. The air provides this heat which causes a reduction in air temperature. This is known as the heat of vaporization and is the same phenomenon we experience when we sweat and feel cooler when a breeze passes by.

Figure 4

Evaporative cooling in high efficiency cooling module

1. Cool Ambient Air
2. Hot IT Return Air
3. Cool IT Supply Air
4. Hot Ambient Exhaust Air
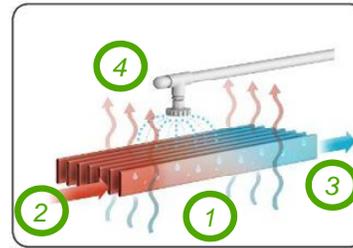
Impact of IT operating environment on economization

In most data centers today, the average IT inlet temperatures ranges from 65-70°F (18-21°C). Many data center operators are very conservative with what they define to be "acceptable" temperature and humidity envelopes for their IT space, because they believe it is necessary to ensure reliable operation and to avoid pre-mature failures of their IT equipment. In contrast to this, ASHRAE TC9.9 recently released its "2011 Thermal Guidelines for Data Processing Environments"[4], which recommends a wider operating environment for temperature and humidity, and IT vendors are specifying even wider acceptable operating windows. **Figure 5** provides a comparison of the original ASHRAE recommended envelope, the new ASHRAE recommended and allowable limits, and typical vendor specifications today.

Figure 5

The expanding operating environment
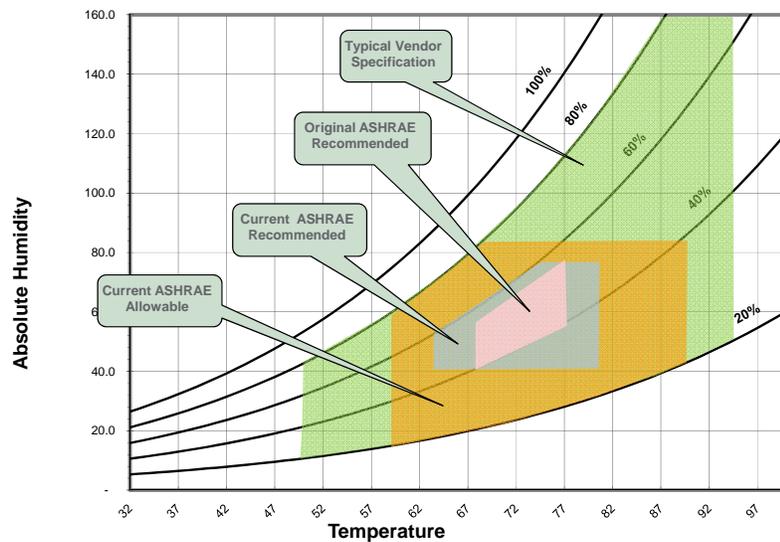
The horizontal axis is the temperature range and the vertical axis is the humidity range. **The bigger the window of acceptable conditions defined for the IT equipment in a data center, the greater the number of hours the cooling system can operate in economizer mode.**

Continuing with the same assumptions from the **Figure 3** comparison, consider the effect of IT inlet temperature on the traditional chiller plant design with a plate and frame heat exchanger. **Table 1** illustrates the increase in hours achieved by increasing the IT inlet temperature to the ASHRAE recommended limit.

---

[4] http://tc99.ashraetcs.org/documents/ASHRAE%20Whitepaper%20-%202011%20Thermal%20Guidelines%20for%20Data%20Processing%20Environments.pdf , accessed on June 22, 2011

Table 1

Impact of increasing IT inlet temperature traditional plate and frame heat exchanger (3 heat exchanges)

IT inlet temperature	Maximum outdoor air temperature	Full economizer hours	% of year on full economizer
70 F (21.1C)	DB: 43.7F (6.5C) MCWB: 39.6F (4.2C)	2,419	28%
80.6 F (27C)	DB: 56.7F (13.7C) MCWB: 51.2F (10.6C)	4,070	47%

Table 2 illustrates the additional hours gained for the high efficiency self contained cooling module (single heat exchange) when the ASHRAE recommended temperature limit is set. The number of hours on economizer mode operation is increased dramatically, to 72% of the year. **As data center operators become more comfortable moving toward the wider environmental envelopes, economizer mode operation can become the primary operating mode rather than the "backup" mode.** Note, as the window of operating temperature and humidity continues to widen, direct air systems (zero heat exchanges) have the opportunity to further increase economizer mode hours.

Table 2

Impact of increasing IT inlet temperature on self contained system (1 heat exchange

IT inlet temperature	Maximum outdoor air temperature	Full economizer hours	% of year on full economizer
70 F (21.1C)	DB: 59.3F (15.2C) MCWB: 53.5F (12C)	4,394	50%
80.6 F (27C)	DB: 72F (22.2C) MCWB: 63,9F (17.7C)	6,289	72%

Efficient cooling components

Variable frequency drives (VFD) on cooling components (i.e. fans, pumps, compressors) and electronically commutated (EC) fans save significant energy. Many data centers today use components in their cooling design that lack VFDs in components including chillers, air handler fans, heat rejection pumps and chilled water pumps. Consider the energy waste from a data center that is 50% loaded, and has air handlers running at 100% (max fan speed). These inefficiencies become even more dramatic when 2N redundancy is a requirement. VFDs address this energy waste by reducing the fixed losses, so the system does not consume as much power at lighter loads.

Integrated controls programmed at the factory

Upon purchasing a hybrid electric vehicle, the expectation is that it will transition smoothly and efficiently between electric and gasoline modes like clockwork, no matter what. This is a common and undisputed expectation due in large part to the standardization of automobiles. This same expectation is possible for standardized, self-contained, economizer-based data center cooling systems. It is only through this level of standardization that an economizer-based cooling system will operate efficiently and predictably in all modes of operation as climate and settings change. Specifically, the controls and management software must be standardized, pre-engineered, programmed at the factory, and wholly integrated into a self-contained cooling system.

Controls in traditional designs, on the other hand, are generally engineered onsite. One-time engineering of the controls and management scheme generally result in controls that:

- are unique and cannot be replicated
- aren't optimized for energy consumption
- aren't fully tested
- don't have fully documented system operation
- are inflexible to data center changes
- require manual inputs and monitoring

This is why it is extremely difficult to build unique economizer-based cooling systems and controls that operate efficiently and predictably in all climates. Some examples of complexities with the controls of a chiller / cooling tower / plate and frame heat exchanger design include:

- Determining the ideal operational points of all the individual components that produce the lowest overall system energy use under any given condition
- Determining all the ways the system can fail and accounting for those failure modes in the custom controls for fault-tolerant operation (e.g. a communications architecture that is tolerant to a severed network cable)
- Determining, controlling, and integrating the disparate non-standard components (pumps, specialized valves, and variable frequency drives (VFDs)) that move the water over the tower's "fill"
- Determining and controlling the fans which may be single speed, multi speed, or variable speed
- Determining the sequence of operation in the event of a step load – in economizer mode, will the chiller need to quickly come online until the chilled water temperature reaches steady state?
- Determining if a chilled water storage tank is required in order to provide the chiller some time to reach stable operation during state changes
- Integrating the chillers with physical features (bypass loops, etc) that allow chillers to sufficiently "warm" the condenser water during the transition from economizer to DX cooling operation (if condenser water is too cold, the chiller will not turn on)
- Controlling basin heaters, or multiple stages of basin heaters, which may be required to prevent freezing when used for winter-time free cooling
- Controlling the integrated valves within the piping architecture, pumps that supply the tower water, and the heat exchangers and chillers that depend on the tower for proper operation

Flexibility / agility improvement

Data centers often demand very flexible architectures in order to accommodate changes in IT requirements while minimizing capital and operating expenses. The cooling approach presented in this paper meets these needs with the following performance attributes:

- Standardized, self-contained design
- Modular design for scaling capacity
- Minimized IT space footprint

Standardized, self-contained design

A standardized and pre-engineered cooling system that is delivered in pre-packaged modules, such as skids, containers, or kits ensures manufacturing in a controlled environment and simplifies shipping and installation. An example of how installation can be simplified is

the use of quick connects in the design to allow for easy hookups to main water supply for the evaporative cooling. White Paper 163, *Containerized Power and Cooling Modules for Data Centers*, discusses the time, upfront cost, and maintenance cost savings, as well as the flexibility, and reliability benefits of standard containerized designs.

Traditional data center cooling infrastructure, on the other hand, can be very complex in the number of components, and how they are installed, maintained, and managed. Installation requires extensive piping, ducting, insulation, and the connection of multiple sub-systems (pumps, chillers, cooling towers, etc) at the site. **Figure 6** illustrates an example of such a design. These many components are often sourced from different vendors, and are custom integrated on the site for the particular installation. This typically means it's more expensive, more time consuming, and more difficult to expand. In addition, they have a higher likelihood of failure and emergency maintenance, as well as a blurred line of responsibility when a failure does occur.
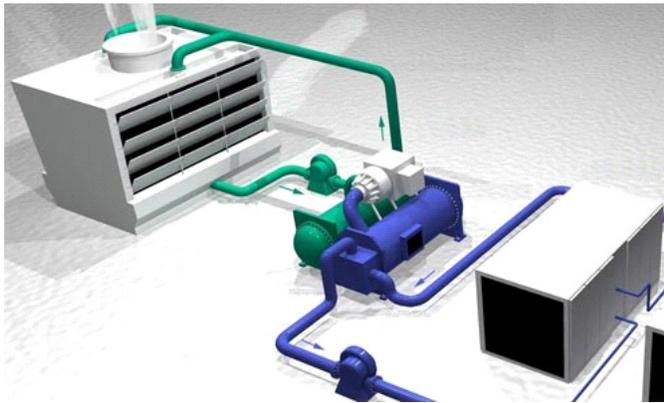
Figure 6

Example of complexity of cooling designs

Scalability

As data centers expand, it is important that its cooling architecture be able to grow with it, rather than overbuild upfront for a worse-case unknown final data center load. Being able to deploy capacity over time helps manage operating expenses and capital expenses. A modular design allows for the capacity to scale as needed, without downtime to the IT load or to the installed system.

In traditional cooling designs, however, cooling components such as chillers and cooling towers are generally sized for the maximum final load of a data center, due to the reliability risk and complexity in scaling such components in production cooling loops. This results in significant overbuilding since load growth is generally very uncertain. Overbuilding this infrastructure means unnecessary capital expenses and operating expenses (to install and maintain more capacity than needed) and decreased efficiency.

Minimize IT space footprint

A self-contained cooling system placed outside the building perimeter or on the roof means more space is available in the IT room for value-added IT equipment. Furthermore, when compared to the total footprint of all the components in a chilled water / cooling tower system, a self-contained cooling system has a smaller footprint. An additional benefit of keeping the cooling system outside of the IT space is that less personnel will need to access the IT space (i.e. for maintenance activities and upgrades / installations), reducing the risk of downtime from human error.

In a typical data center, 10-20% of the white space is consumed by physical infrastructure components including air handlers / air conditioners, humidifiers, UPSs, power distribution units and the required service clearances. This is space that cannot be used for the value-added IT equipment. In some parts of the world where real estate is at a premium, this is a significant limitation of data center design.

Reliability and availability improvement

The primary goal of most data center managers is to ensure that the critical IT loads remain operational at all times. A cooling system that addresses reliability and availability needs of today's data centers must:

- be fault tolerant and maintainable without going off-line
- isolate indoor air from outdoor air for a controlled environment
- minimize its dependence on utility water
- address the environmental concerns over chemicals associated with some refrigerant or water-based systems
- provide predictable airflow through containment of supply and return air streams

Maintainability

Maintaining IT operations while servicing the cooling plant is critical to achieving reliability and availability goals. Many cooling systems require a complete system shutdown for certain maintenance activities. This means, in order to have concurrent maintenance, a complete 2N system is required, which is very costly. For example, with a chilled water design, the data center would need two independent chillers so that one could continue to operate and cool the data center while the other was being serviced. In some cases, an N+1 design may meet the concurrent maintenance requirements. A self-contained system designed with device redundancy avoids this additional expense while still achieving concurrent maintainability.

Another maintenance consideration is the risk of downtime from human error during the maintenance activity. In chiller plant designs, air handlers are located inside the IT space; therefore maintenance on the air handlers means personnel are working in a live IT operating environment. A system completely located outside reduces downtime risks because the service personnel are not performing their work inside the IT space.

Controlled environment

A system with an air-to-air heat exchanger and evaporative cooling provides significant energy savings over typical cooling approaches, while still ensuring complete separation of indoor and outdoor air. This is important for those data center managers concerned about contaminants, clogged air-filters, or swings in temperature and humidity that could increases the downtime risk of their IT equipment.

Minimize dependence on utility water

A system with a lower dependence on utility water throughout the year is less likely to experience a failure due to the loss of utility water. With a chilled water / cooling tower cooling design, the data center's operation is dependent on the delivery of utility water. Loss of the utility water would mean the cooling tower is left without makeup water, which the system is dependent on 8760 hours of the year. Cooling towers consume approximately 40

gallons per minute / 1,000 tons of cooling capacity (151.4 liters per minute)[5]. Improved architectures, such as the self-contained system discussed in this paper, do use water for evaporative assist, but to a much lesser extent since it only uses the evaporative assist process during the hotter periods of the year. The probability that the loss of the utility water would occur at the same time as the operation of the evaporative assist is much lower.

Environmentally-friendly

As part of their "green" company initiatives, some data center managers are looking for options that address the environmental concerns over chemicals associated with some refrigerant or water-based systems.

A cooling system with a chemical-free water treatment system eliminates all contaminants in the water including potential bio-threats. A common type of chemical-free system sends electrical pulses through the water to change polarity of mineral contaminants which causes them to clump together and precipitate out into powder form and then get flushed out of the sump. Micro-organisms get encapsulated by this clumping action and, by passing through electrical pulses, their cell walls are damaged through electroporation. This causes them to spend their short life cycle trying to repair themselves rather than reproducing and posing a threat to the water system. Such a system eliminates the costs of chemicals and special maintenance of chemical treatment, and addresses the environmental concerns. In addition, the blow down water from such a system can be reused for gray water usage at the facility, conserving water consumption.

Predictable airflow performance

Air containment, to separate hot return air from cold supply air is crucial to efficient cooling. Without a form of air containment, either hot spots are likely – something data center managers try at all costs to avoid – or significant over-provisioning of the coolers occurs, which means a significant increase in energy consumption and overall costs. White Paper 135, *Hot-Aisle vs. Cold-Aisle Containment for Data Centers*, discusses the challenges of air mixing, and provides recommendations for effective containment of the air in new data centers.
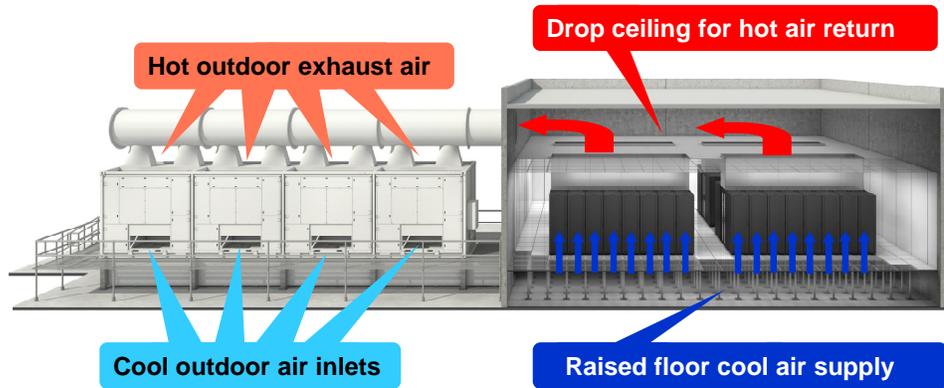
The IT space can be a raised floor environment with perforated tiles for air distribution like typical data centers, or air can be distributed with air diffusers at row ends to deliver the air to the IT on cement slabs. Hot air from the servers is controlled through ducting connected to the racks. The hot air rises to the ceiling plenum and is fed into the return ducting of the cooler. **Figure 7** illustrates how the supply and return air in a self contained cooling module is ducted into the IT space in a raised floor environment. Regardless of the cooling plant architecture used, separation of hot and cold air is a best practice that should be adopted by all data centers to improve efficiencies and cooling performance.

---

[5] Arthur A. Bell, Jr., *HVAC Equations, Data, and Rules of Thumb* (New York: McGraw-Hill, 2000), p. 243

Figure 7

Air distribution of indirect air cooling plant

Hot outdoor exhaust air

Drop ceiling for hot air return

Cool outdoor air inlets

Raised floor cool air supply

Comparison of cooling architectures

Data center designers and managers face the difficult decision of choosing between numerous cooling architectures. TradeOff Tool 11, *Cooling Economizer Mode PUE Calculator*, helps quantify this decision, and illustrates which architecture(s) have the optimal PUE, energy cost, and carbon emissions for their data center location and IT operating environment. **Figure 8** illustrates the inputs and outputs of this tool.

Figure 8

TradeOff Tool Calculator to help assess performance of various cooling approaches
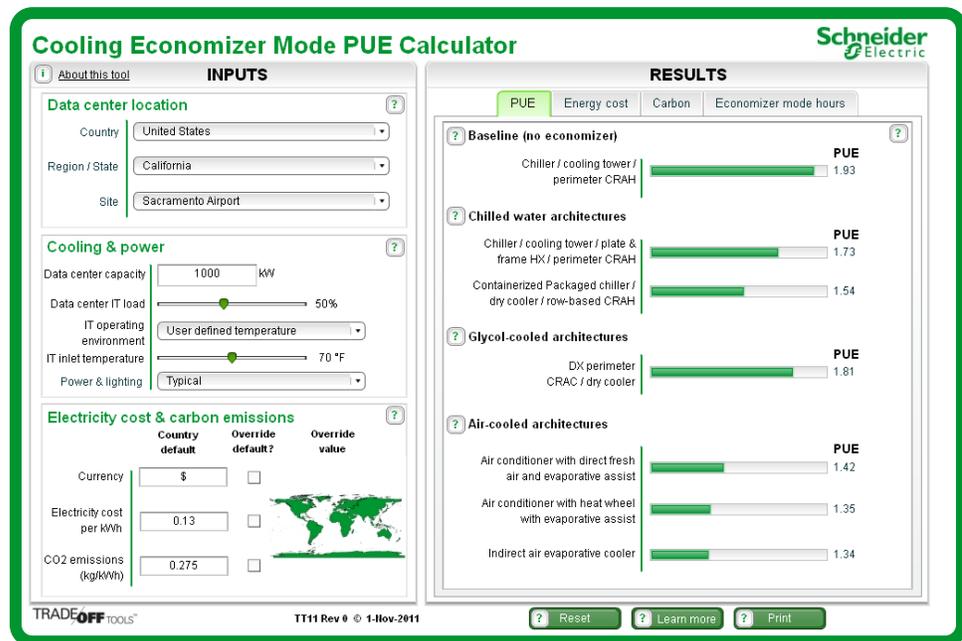
Table 3 provides a comparison of two architectures – a traditional chiller plant design (defined in the text box below) with a plate and frame heat exchanger, and a self-contained cooling system (as discussed in the earlier parts of this paper). The self-contained cooler provides significant benefits over the traditional approach, as noted by the highlighted cells of **Table 3**.

> **> Traditional cooling method**
> A traditional cooling method is defined as having the following attributes:
>
> - CRAC/CRAH units are located in the IT room
> - Air is distributed under raised floor via vented tiles
> - Outdoor heat rejection is via cooling tower
> - Components are installed upfront for maximum projected cooling capacity needed
> - System has minimal economizer mode operation
> - Cooling components are from various manufacturers and are integrated for a project
> - Controls are created for the project
> - Management software is customized for the project

Table 3

Comparison of cooling performance

Design characteristic	Self-contained Indirect Evaporative Cooler	Traditional chilled water plant
Primary mode of operation	Economization modes (air-to-air heat exchanger and evaporative cooling) with DX coolers as backup	Chiller operation with plate and frame heat exchanger as backup
Controls and management software	Standardized, pre-integrated controls ensures optimal operation mode at all times; few devices to control	Many devices to control; complex custom controls often result in a cooling plant not in optimal mode of operation
Form factor	Self-contained in one unit that is fully integrated	Chillers, pumps, cooling towers, and piping are disparate parts that are assembled and integrated in the field.
IT space footprint	Zero IT space footprint; sits outside the data center	Consumes approximately 30 sq m for every 100 kW of IT load, or approximately 5% of computer room space
Ability to retrofit	Not logical to retrofit into existing facilities; only cost effective for new facilities	Practical if space is available; requires running additional pipes
Energy use	Operates in economizer mode > 50%* of year; One heat exchange means economizer mode can run at higher outdoor temperatures * Based on assumptions of Figure 3	Operates in economizer mode approximately 25%* of year; Primary mode of operation is full mechanical cooling; Three points of heat exchange means greater temperature difference required between IT inlet temperature and outdoor temperature * Based on assumptions of Figure 3
Dependence on water	Lower probability of losing water at the same time evaporative assist is required	Loss of utility water is critical – cooling tower depends on makeup water 8760 hours of the year
Controlled environment	Outside air contaminants are isolated from IT intakes reduces risk of downtime	Outside air contaminants are isolated from IT intakes reduces risk of downtime
Upfront cost	$2.4 / watt for entire system	$3.0 / watt for entire system

Conclusion

Today's data center managers are facing increased financial and regulatory pressure to improve the efficiency of their data centers. In order to achieve the aggressive PUE targets being set by management, data center managers must adopt the cooling philosophy that the primary mode of operation is on economizer, and the mechanical system is the back-up to the economizer when needed. For a significant number of climates across the globe, an indirect evaporative cooling system with air-to-air heat exchange is the most effective way to achieve this, without exposing the IT space to the outside air contaminants and conditions directly.

In addition, data center managers must look for a cooling architecture that can adapt effectively to varying IT loads, can be scaled quickly as capacity is needed, and is standardized and pre-engineered with integrated controls for optimal operation. Along with best practice airflow management and a wider operating window for IT temperature, cooling capex and opex can be reduced substantially.

Tools such as Schneider Electric's Cooling Economizer Mode PUE Calculator can help identify the optimal cooling architecture for a specific geographic location and IT load characteristics.

✎ About the author

Wendy Torell is a Senior Research Analyst at Schneider Electric's Data Center Science Center. She consults with clients on availability science approaches and design practices to optimize the availability of their data center environments. She received her Bachelor's of Mechanical Engineering degree from Union College in Schenectady, NY and her MBA from University of Rhode Island. Wendy is an ASQ Certified Reliability Engineer.

Resources

Economizer Modes of Data Center Cooling Systems
White Paper 132

Specification of Modular Data Center Architecture
White Paper 160

Containerized Power and Cooling Modules for Data Centers
White Paper 163

Hot-Aisle vs. Cold-Aisle Containment for Data Centers
White Paper 135

Browse all white papers
whitepapers.apc.com

Cooling Economizer Mode PUE Calculator
TradeOff Tool 11

Browse all TradeOff Tools™
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm