

Beyond Chiller Efficiency: Optimizing Total Data Center Cooling System Performance

In the complex world of data center cooling, optimizing for peak efficiency is paramount. However, a narrow focus on individual component performance can often obscure a more significant opportunity: optimizing the entire system as a cohesive whole. This analysis reveals a counter-intuitive but powerful strategy for water-cooled plants. By strategically raising the condenser leaving-water temperature, we can accept a minor decrease in chiller efficiency to unlock substantial gains across the entire cooling system. This holistic approach can lead to significant benefits, including reduced initial costs, a smaller equipment footprint, and lower annualized energy consumption, demonstrating that the greatest efficiencies are found when we evaluate performance at the plant level, not just the component level.

Data center thermal infrastructure designs utilizing water-cooled chillers with dry fluid coolers offer an opportunity to optimize overall system efficiency by operating at elevated condenser leaving-water temperatures, provided equipment operating limits are maintained. While increasing condenser leaving-water temperature raises chiller lift and can reduce chiller efficiency at the unit level, the resulting increase in dry fluid cooler heat rejection capability can help improve total system performance. The higher condenser water temperature increases the temperature differential between the condenser loop and ambient air, which can either reduce the number of dry fluid coolers required or allow the installed heat rejection equipment to operate at reduced fan and pump speeds. These effects can help lower first cost, reduce plant footprint, and/or decrease annualized cooling energy consumption (improved PUE). This paper examines these tradeoffs for hyperscale and colocation data center applications and demonstrates that, in many cases, shifting heat rejection burden from less efficient auxiliary components to the chiller may help improve overall cooling system efficiency when evaluated at the system level rather than the component level.

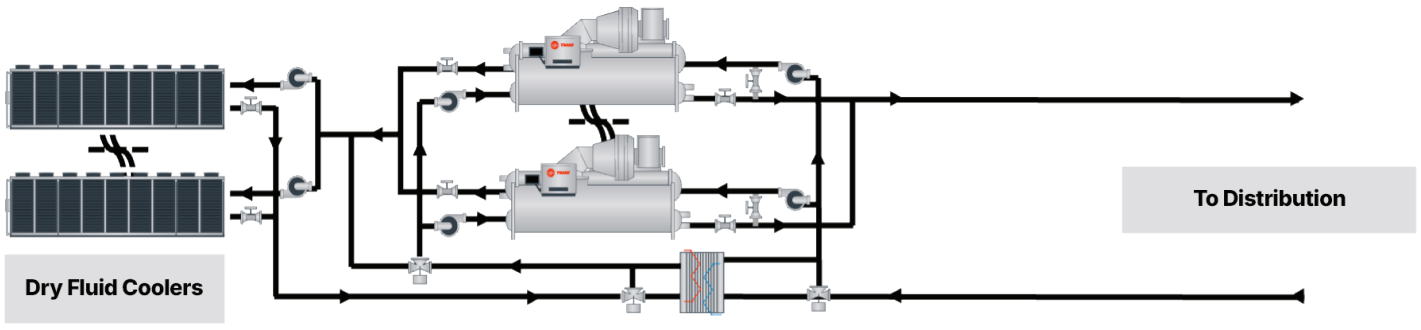
Data centers utilizing water cooled chillers in conjunction with dry fluid coolers can operate over a flexible condenser water temperature range, subject to the equipment's operational ranges. Desired condenser water temperatures can be maintained either by selecting dry fluid coolers rated for higher approach or entering fluid temperatures, or by leveraging the dry coolers' ability to modulate fluid flow and/or fan speed. This operational flexibility allows engineers to optimize total system performance, system energy use, acoustical impact, and overall heat rejection efficiency. Operating at a higher condensing water temperature can reduce the number of dry fluid coolers required, lowering first cost and footprint or, with the original cooler count maintained, enable more efficient operation and/or lower sound levels. For hyperscale and colocation owners and operators, condenser water temperature is not just a design setting—it is a business lever that can help influence capital cost, operating efficiency, site footprint, and long-term asset value.

System Layout

Consider the system shown in Figure 1. The cooling distribution loop serves a typical IT load, with chilled water provided by water cooled chiller(s). Condenser heat from the chiller is rejected to the atmosphere via dry fluid coolers. A waterside economizer heat exchanger is incorporated to enable free cooling whenever outdoor ambient conditions allow.

The evaporator leaving water temperature is established based on the IT load's cooling requirements. The condenser leaving water temperature is selected in coordination with the dry fluid cooler's heat rejection capacity under local design conditions to ensure the chiller operates within its allowable performance envelope.

Figure 1 – Example System Layout

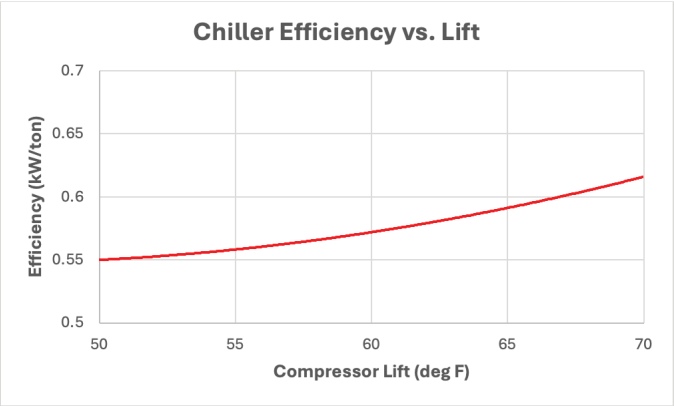


Learn more about Trane® Continuum Rubin DSX Reference Designs.

Impact To Chiller

Assuming a constant evaporator leaving water temperature, increasing the condenser leaving water temperature raises the chiller's lift. Higher lift drives the chiller to a less efficient operating point (see Figure 2); however, in most applications, the chiller is still the most efficient component in the thermal management system. By reallocating some of the thermal rejection workload from a less efficient device to the more efficient chiller, the overall system efficiency can be improved. Ultimately, optimizing around total cooling system performance—not individual equipment efficiency—can help reduce energy use and improve operating economics at scale.

Figure 2 – Example Chiller Efficiency Curve

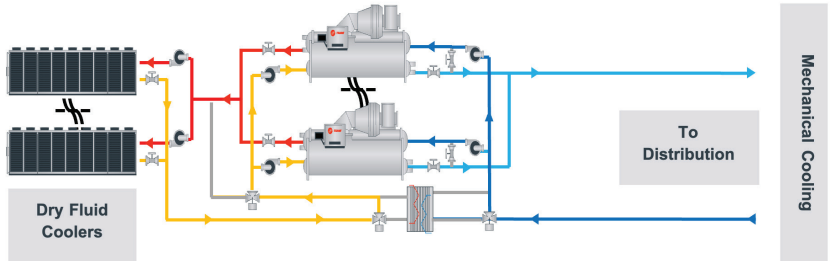


Impact on Dry Fluid Coolers

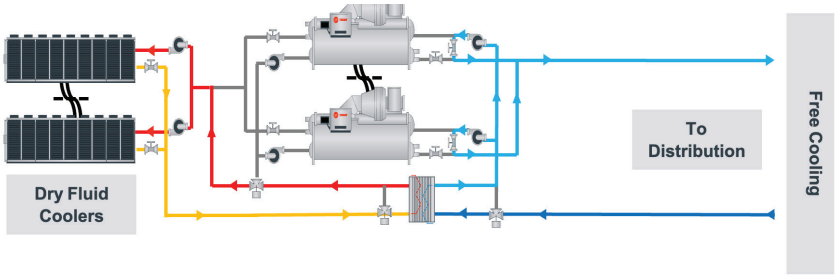
Dry fluid coolers in this configuration serve two functions. During mechanical cooling, they reject heat from the chiller's condenser water loop. During free cooling operation, the same dry coolers reject heat from the facility water loop. A hybrid mode exists that provides unloading to the mechanical cooling when full free cooling capacity is not available.

In general terms, the cooling modes can be defined as:

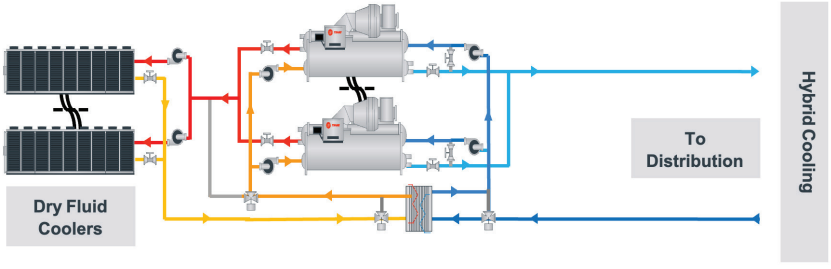
Mechanical Cooling =
Ambient temperature >
Fluid return temperature



Full Free Cooling =
Ambient temperature <
Required supply temperature



Hybrid Free Cooling =
Fluid return temperature >
Ambient temperature >
Required supply temperature



[Learn more about Trane® Continuum Rubin DSX Reference Designs.]

Increasing the condenser leaving water temperature increases the temperature differential between the condenser fluid (heat source) and the ambient air (heat sink). A larger temperature differential allows each dry cooler to deliver more heat rejection capacity than it would with a smaller differential. As a result, the system may require fewer dry coolers overall, or the existing dry coolers can operate at more favorable—and typically more efficient—conditions. In many designs, the actual benefit is a blend of both outcomes. For hyperscale and colocation owners/operators, this means condenser leaving-water temperature is not just a mechanical setting, but a strategic lever that can help reduce infrastructure requirements, lower auxiliary energy use, and improve total cooling plant economics.

When comparing a baseline condenser water temperature and its associated dry fluid cooler selection to elevated condenser water temperature operation, the advantages and disadvantages across typical scenarios include the following:

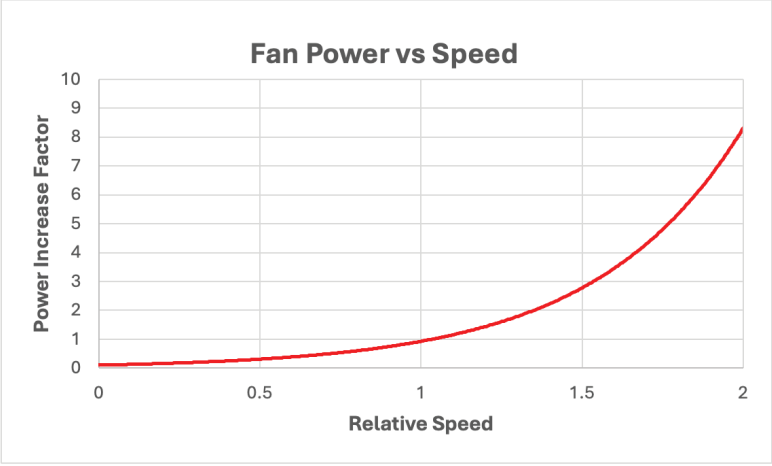
Scenario 1 – Fewer Dry Coolers/Pumps

Increasing the condenser water temperature raises the heat rejection capacity of each dry fluid cooler. With higher dry fluid cooler unit capacity, fewer dry fluid coolers may be required to meet the design load. This reduces first cost, decreases the required installation footprint, and lowers total fan power due to fewer fans operating. Under this configuration, the ambient air temperature at which free cooling becomes viable is lowered, enabling fewer hours of free cooling economizer operation.

Scenario 2– Maintain Quantity of Dry Coolers

If the original number of dry fluid coolers is retained, the ambient air temperature at which free cooling becomes viable remains unchanged. This preserves the original number of free cooling economizer hours and maintains the same physical footprint. However, operating at a higher condenser water temperature allows the dry coolers to run at reduced fan speeds, lowering total fan energy in accordance with fan affinity laws (Fan power affinity laws state that the power changes with the cube of the relative speed change), resulting in an improved PUE for the data center. An added benefit of the reduced fan speed is lower acoustic impact to the surrounding area.

Figure 3 – Affinity Law Power Curve $P_1/P_2=(n_1/n_2)^3$



Impact To Overall System

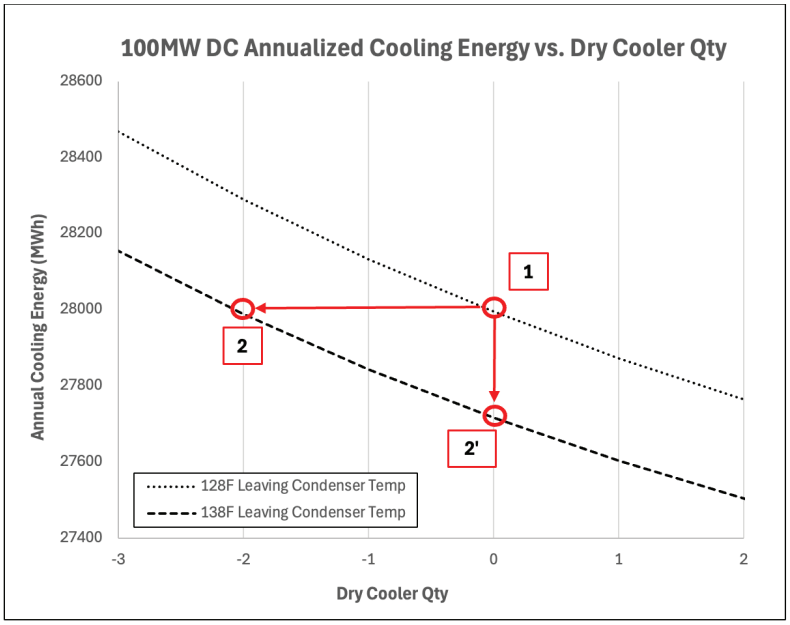
Consider a theoretical 100 MW data center served by the previously described water-cooled chiller / dry fluid cooler configuration. The system is designed with a 75°F supply temperature and a 104°F design ambient condition. Two condenser-water leaving temperature setpoints were evaluated: 128°F and 138°F.

Figure 4 illustrates the two operating strategies described previously. Point 1 represents the baseline case, with a 128°F condenser-water leaving temperature and the corresponding dry fluid cooler quantity required to satisfy the design basis. Increasing the condenser-water leaving temperature to 138°F shifts operation to Point 2, where the higher heat-rejection temperature enables a reduction in the number of dry fluid coolers required, consistent with Scenario 1. This approach reduces both first cost and installation footprint.

Alternatively, shifting from Point 1 to Point 2′ maintains the original dry fluid cooler quantity while reducing the plant's annual cooling energy consumption (improved PUE), consistent with Scenario 2.

Note: *While these two scenarios highlight discrete system design directions, they can be combined to arbitrate between the benefits described in each individual scenario.*

Figure 4 – Higher Condensing Water Temperature Benefits



Conclusion

For data centers utilizing water-cooled chillers and dry fluid coolers, condenser leaving water temperature should be evaluated as a system optimization variable rather than a fixed design constraint. Although increasing condenser leaving water temperature moves the chiller to a higher lift and less efficient operating point, the corresponding improvement in dry fluid cooler performance can create meaningful benefits at the system level. Depending on project objectives, these benefits may include a reduction in dry fluid cooler and pump quantity, lower installed cost, reduced footprint, lower fan energy, reduced acoustical impact, or reduced annualized cooling energy through improved operation of the existing heat-rejection infrastructure. For hyperscale and colocation owners and operators, this means condenser-improved PUE resulting in less power consumption and/or more compute power available for token generation.



The analysis suggests that the preferred condenser water temperature strategy depends on the relative value placed on first cost, footprint, free cooling economizer hours, and lifecycle energy performance. In applications where the chiller remains the most efficient heat rejection device in the system, it may be advantageous to accept a moderate chiller efficiency penalty in exchange for larger reductions in auxiliary power. As a result, elevated condenser leaving water temperatures can represent a practical and effective design strategy for data center cooling systems, particularly when assessed using total system performance metrics.

About the Author

Daryl Keys is a Lead Engineer in the Integrated Solutions Team, bringing over 27 years of expertise in HVAC and hydronic systems. His current role focuses on hydronic systems development for data center and commercial building applications.



To learn more about our innovations in thermal management systems for data centers, visit Trane.com/DataCenters or contact your Trane Account Manager. For more information on Trane Continuum Reference Designs, visit Data Center Reference Designs | Trane Commercial HVAC.

All trademarks referenced in this document are the trademarks of their respective owners.

Trane – by Trane Technologies (NYSE: TT), a global climate innovator – creates comfortable, energy efficient indoor environments through a broad portfolio of heating, ventilating and air conditioning systems and controls, services, parts and supply. For more information, please visit *trane.com or tranetechnologies.com.*

© 2026 Trane. All Rights Reserved | DC-WPR008A-EN | 10/16/2026

This is for informational purposes only. Trane believes the facts and suggestions presented here to be accurate. However, final design and application decisions are your responsibility and will affect actual financial and energy efficiency results. Trane disclaims any responsibility for actions taken on the material presented.