DOI:10.20357/B7RC8M

Lawrence Berkeley National Laboratory

# The Installation of Direct Water-Cooling Systems to Reduce Cooling Energy Requirements for High-Performance Computing Centers

Shankar Earni[1], Steve Greenberg[1], and Walker Johnson[2]

[1]Lawrence Berkeley National Laboratory, [2]kW Engineering

Energy Technologies Area
August 2020

Disclaimer:

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

# The Installation of Direct Water-Cooling Systems to Reduce Cooling Energy Requirements for High-Performance Computing Centers

*Shankar Earni and Steve Greenberg, Lawrence Berkeley National Laboratory*
*Walker Johnson, kW Engineering*

**ABSTRACT**

A large cluster of High-Performance Computing (HPC) equipment at the Lawrence Livermore National Laboratory in California was retrofitted with an Asetek cooling system. The Asetek system is a hybrid scheme with water-cooled cold plates on high-heat-producing components in the information technology (IT) equipment, and with the remainder of the heat being removed by conventional air-cooling systems. In order to determine energy savings of the Asetek system, data were gathered and analyzed two ways: using top-down statistical models, and bottom-up engineering models. The cluster, "Cabernet", rejected its heat into a facilities cooling water loop which in turn rejected the heat into the same chilled water system serving the computer-room air handlers (CRAHs) that provided the air-based cooling for the room. Because the "before" and "after" cases both reject their heat into the chilled water system, the only savings is due to reduction in CRAH fan power. The top-down analysis showed a 4% overall energy savings for the data center (power usage effectiveness (PUE) —the ratio of total data center energy to IT energy— dropped from 1.60 to 1.53, lower is better); the bottom-up analysis showed a 3% overall energy savings (PUE from 1.70 to 1.66) and an 11% savings for the Cabernet system by itself (partial PUE of 1.51). Greater savings, on the order of 15-20%, would be possible if the chilled water system was not used for rejecting the heat from the Asetek system. About 37% of the heat from the Cab system was rejected to the cooling water, lower than at other installations.

## Introduction

Data centers are an essential and important part of daily life where their use is ubiquitous from a simple web search to running complex machine-learning algorithms. Their usage is on the upward trend and will likely go even higher considering various emerging applications like blockchain and autonomous driving. As a result the energy use to run them and the supporting infrastructure would likely increase proportionally, though there is a balancing trend toward more efficient processing and more efficient infrastructure. According to some estimates, data centers account for around 1% of worldwide electricity use. Data centers in the United States consumed an estimated 70 billion kWh—1.8% of total U.S. electricity consumption (Shehabi et al. 2016).

A big portion of this electricity consumption is going towards removing the heat that's generated in the processing of various transactions and computations. All of the electrical input power to the IT equipment is converted into heat that must be removed by appropriate cooling. The type of cooling system depends on the size and type of the data center, along with its operating characteristics, including the outside weather conditions. In general air-cooled systems (Capozzoli and Primiceri, 2015) are the most prevalent form of cooling system and these systems have evolved to cope with the advances in information technology equipment (ITE). These

advances are moving towards localized cooling units thereby bringing them closer to the data center to support their rising power densities. Cooling in data centers can broadly be classified based on the proximity for transferring heat from the source ranging from farthest (like computer room air handlers (CRAHs)) to closest (immersion cooling where the heat transfer takes place on the surface of the hot electronic components).

Data center cooling systems can be room-based like CRAHs that use chilled water to control temperature and humidity in the server room. The heat rejection can be done either through evaporating water (e.g. with a cooling tower) or without evaporation (e.g. a dry cooler). . Rows of ITE racks are conventionally arranged with alternating cold (IT inlet air) aisles and hot (IT discharge air aisles. Thus the ITE is arranged with the fronts facing each other across the cold aisle(s) and the backs facing each other across the hot aisle(s). Some of the cooling systems are more localized than CRAHs, like in-row coolers that are placed in the middle of a row of racks (Patankar 2010) and draw in hot air from its rear (hot-aisle) and exhaust cold air into the cold aisle after internally cooling the air. There are other systems like overhead-cooling units that draw in hot air from the hot aisle, cool it, and blow cold air downwards into the cold aisle. Other cooling systems are even more targeted and closer to the racks where the heat is generated. Rack-based systems can be based on a closed design, where the servers and the terminal cooling equipment are accommodated within the closed rack envelope and thus the airflow is completely isolated from the rest of the data center. These are known as in-rack coolers.
Rack-based systems can also be based on an open design, which is typically characterized by a rear-door heat exchanger that is mounted as the rear door on a rack, which shares with a room-based solution the task of cooling the ITE. The rear door can be active, if it has dedicated fans to control the airflow through the back-door heat exchanger. Rear-door systems are considered to be passive if the server fans drive the airflow through the rear door.

Other cooling systems include those where liquid flows through a server and makes contact through micro-channel flow and cold-plate heat exchangers in direct contact with some components, such as central processing units (CPUs) and dual in-line memory modules (DIMMs). A study of such liquid-cooling systems (Coles and Greenberg 2014), that tested a range of configurations, found that the typical percentage of heat captured by the liquid was around 50-60%, although it depends on the environmental conditions and IT load. Liquid cooling provides a great option to reduce data center cooling energy requirements because the heat capacity of liquids is orders of magnitude greater than that of air and once heat has been transferred to a liquid, it can be removed from the data center efficiently.

This work documents and validates the performance of a liquid-cooling system serving the CPUs installed at one of the high-performance computing clusters. The purpose of the retrofit was to demonstrate the energy-saving potential of this technology.

## Building and Project Description

The data center, at the Lawrence Livermore National Laboratory (LLNL) in California, was equipped with conditioned air through a raised floor with a 3-foot underfloor plenum, and with open air return. The data center is conditioned primarily by 37 CRAH units with chilled water coils and fans equipped with variable frequency drives (VFDs). The chilled water coils in the CRAH units have 3-way valves, controlled to a constant supply air temperature of 64° F by a building automation system (BAS). The fan VFDs are controlled based on return air temperature. There are three centrifugal, water-cooled chillers that provide chilled water to the CRAH units.

The chilled water system has a primary-only pumping configuration, with four constant speed chilled water pumps (50HP each). The chillers are served by the campus condenser water loop. There are no condenser water pumps in the building. Humidification and dehumidification in the data center is achieved through two make-up air units that supply a total of 15,000 cfm of outside air. Relative humidity is kept within roughly 35-45%. Since the CHWST is 46°F, there is also some incidental dehumidification at the CRAHs. The data center has one high-performance computing cluster, Cabernet (Cab), which has been retrofitted with a liquid-cooled Asetek system. A generic schematic of this system appears in Figure 1.

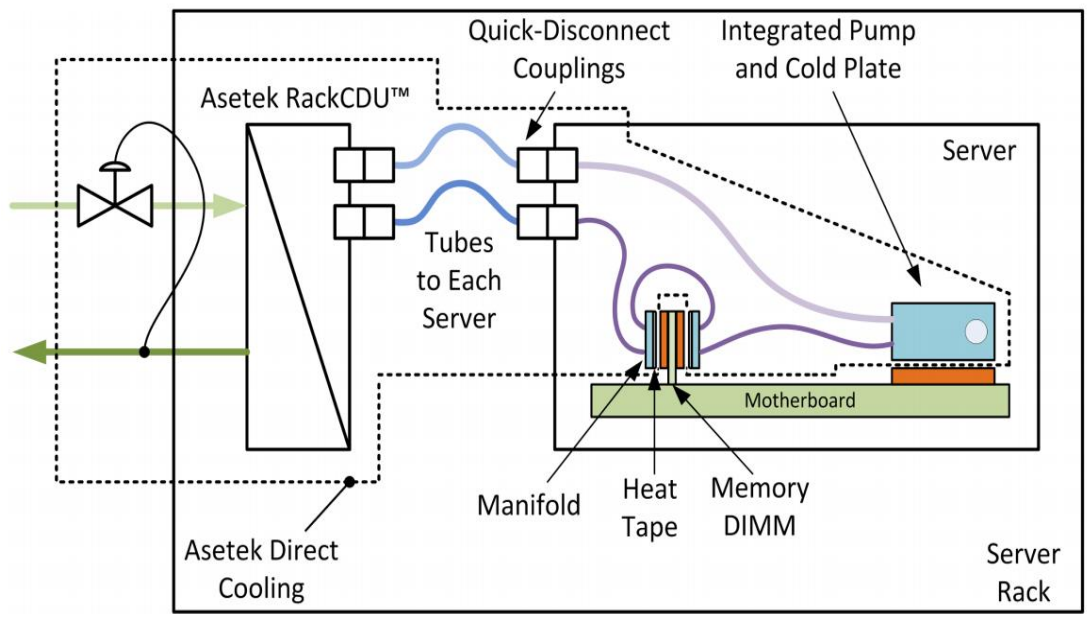


Figure 1: Schematic diagram of liquid cooling system (Coles and Greenberg 2014). Note that in the Cab implementation, only the processors and not the DIMMs were equipped with heat exchangers.

The Asetek system uses direct liquid cooling to cool some components of computer servers (processors in this case) thereby providing part of the necessary cooling; the remainder of the heat is removed by air moved by server fans. The liquid cooling system utilizes two 4W pumps for each server to circulate a glycol and water solution directly through the server. According to Asetek, there are 2464 of these pumps serving Cab. A central cooling distribution unit (CDU) distributes cooling water to an Asetek RackCDU™ at each of 18 racks. The RackCDU™ is a water-to-water heat exchanger with the server loops on the hot side and a secondary loop on the cold side, the latter exchanging heat with a central CDU. The central CDU, with a heat exchanger and pumps to circulate water in the secondary loop, in turn rejects the heat from the secondary loop into the building's chilled water system. Transferring heat from the servers to the chilled water loop bypasses the CRAH units, therefore saving CRAH fan energy, and also saving chiller energy. There are no savings of CHW pump energy as the pumps are constant speed and the CRAH control valves are three-way, resulting in constant flow and pump energy regardless of load.

Also, since the CRAH unit fans are controlled based on return air temperature, the colder air coming back from the servers allows the CRAH fans to slow down. The system installed in

the Cab cluster has one CDU that provides cooling to the 18 racks in the cluster. The energy conservation measure (ECM) is intended to save energy used at the data center by transferring some of the heat from the Cab computer cluster to the building chilled water system directly compared to relying completely on the existing cooling system that uses only air and computer room air handlers (CRAHs). Therefore, the existing CRAH units are expected to "turn down" (reduced fan speed and lower cooling water flow rate) because the cooling load on the CRAH units will be less than during the pre-retrofit condition with the same IT load. The heat not captured by the ECM will be rejected to the outside using the existing CRAHs and chilled water system.

## Data Collection

In order to assess the baseline and post retrofit performance, equipment and operational data was collected from a variety of sources including meters, submeters and spot measurements from equipment control screens. They include Asetek RackCDU heat loads, Cab IT power (kW), remaining data center IT power (kW), chilled water pump (kW), chiller power (kW), CRAH fans speed and power (kW) and fan speed for all CRAH units, along with their supply and return air temperatures. No data were available for the condenser water plant, so we assumed a constant performance of 0.58 kW/ton based on a recommendation from the site personnel.

The facility is equipped with the PI database (a system provided by OSIsoft), the site's overarching system that acts as a backbone for the collection of data from different data sources and devices for record keeping and analysis. Based on the available data, the following data points (Table 1) were available to analyze the performance of the system.

Table 1: Summary of the data variables available.

| Category | System | Variable | Description |
|---|---|---|---|
| IT Load | IT Load 1, incl. Cab | TB1864 (kW, kWh) | Power draw by the IT equipment, partial |
| | IT Load 2 | TB1865 (kW, kWh) | Power draw by the IT equipment, partial |
| | Cab Load | Cab power (kW, kWh) | Power draw by water cooled Cab cluster (included in IT Load 1) |
| Mechanical Load | Chiller Load | TB821 (kW, kWh) | Power draw from all CRAH units. Also feeds some pumps in the chiller room. |
| | CRAH Load | TB1866-(kW, KWh) | Chiller power, also support the office area also supported by TB823. |
| Office Load | Other loads including lighting | TB823-(kW, KWh) | Load from office rooms, lighting and plug loads, assumed that these loads/people support the data center. |
| CDU data (Post retrofit) | Heat captured by the water in each of the RackCDUs | RackCDU Heat Load<br>RackCDU Facility Temperature In<br>RackCDU Facility Temperature Out<br>RackCDU Facility Flow | Individual data from 16 different rack CDUs |
| Other | Weather | Outside air temperature (F) | |
| | Cluster/node data | Internal room, and server inlet temperatures | |

The water-cooled Asetek system was installed, commissioned and became operational in December 2016. The baseline data gathered period is from January 1st 2016 to December 31st 2016. However, the data starting in early July was somewhat anomalous when compared to the first 6 months in terms of shift in the mean and higher level of variability for the mechanical load. Through talking to the installation team, it was found that the installation of the Asetek system started around that time and some of the data might not represent true baseline conditions. Hence we have chosen the first 6 months of 2016 data to develop the baseline. The data that was available from the PI system included variables listed above including their time stamp. The data that was used for analysis is sampled at a 30-minute interval.

The savings for the Asetek retrofit was primarily from reducing load on the CRAH units, allowing for their fans to ramp down and thus saving fan energy. In addition, the reduced fan energy results in less heat being rejected to the chilled water loop, reducing the energy consumption by the chilled water and condenser water plants. A small amount of energy also was added by the Asetek pumps and CDU circulating pump. Only the condenser water plant efficiency is impacted by outdoor air conditions, and since data for the condenser water plant was not available, there was no need (or ability) to develop weather-based calculations. Instead we assumed steady-state operations.

The baseline CRAH fan power was determined by summing the average power consumption from the baseline trends. The same was done for the post-retrofit trends to determine the total CRAH fan savings. Since trends do not indicate other significant changes to the data center during that time, the savings can be attributed to the retrofit. We could not use the chiller power trends directly to calculate chilled water plant savings, because of the trend data discrepancy described above. Instead, we first calculated the baseline chilled water plant performance (in kW/ton) based on the trended chiller power (kW); pump power (kW) and cooling load (IT load plus CRAH fan power). We then multiplied the reduced fan power of the CRAH fans (the cooling load reduction) by the combined chilled water plant and condenser water plant performance to determine the cooling power savings.

**Bottom-Up Model Development and Results**

An energy model was developed by kW Engineering to create energy balances of the whole data center and two different sub-systems. The model is an 8760-hour bin analysis using commercial spreadsheet software and custom formulas based on the ratings of the equipment at the data center. The energy balance shown in Figure 2, below, shows the calculated baseline annual and post retrofit energy use percentages for all systems in the data center. Figure 3 shows the energy balance of the Cab IT load and the associated systems serving it. According to trends, the Asetek system removes 33% of the heat generated by Cab, with the rest of load being removed by the CRAH units. We also used the energy model we developed to calculate the annual PUE for the data center and the annual partial PUE (pPUE) for various sub-systems. The pPUE is the PUE of one individual system, referencing only the IT load that it serves. We included the following:

- Data Center - Baseline and Post-Retrofit PUE
    - CRAHs
    - Chilled Water Plant
    - Condenser Water Plant
    - Total IT Load
- Cab pPUE

    - CRAHs (only the energy associated with Cab)
    - Chilled Water Plant
    - Condenser Water Plant
    - Asetek Pumps (Rack pumps and CDU circulating pump)
    - Cab IT load
- Asetek w/ CHW pPUE
    - Asetek pumps (CDU circulating pump)
    - Chilled water plant (Associated load)
    - Condenser water plant (Associated load)
    - Cab IT load
- Asetek Only pPUE
    - Asetek pumps (CDU circulating pump)
    - Cab IT load

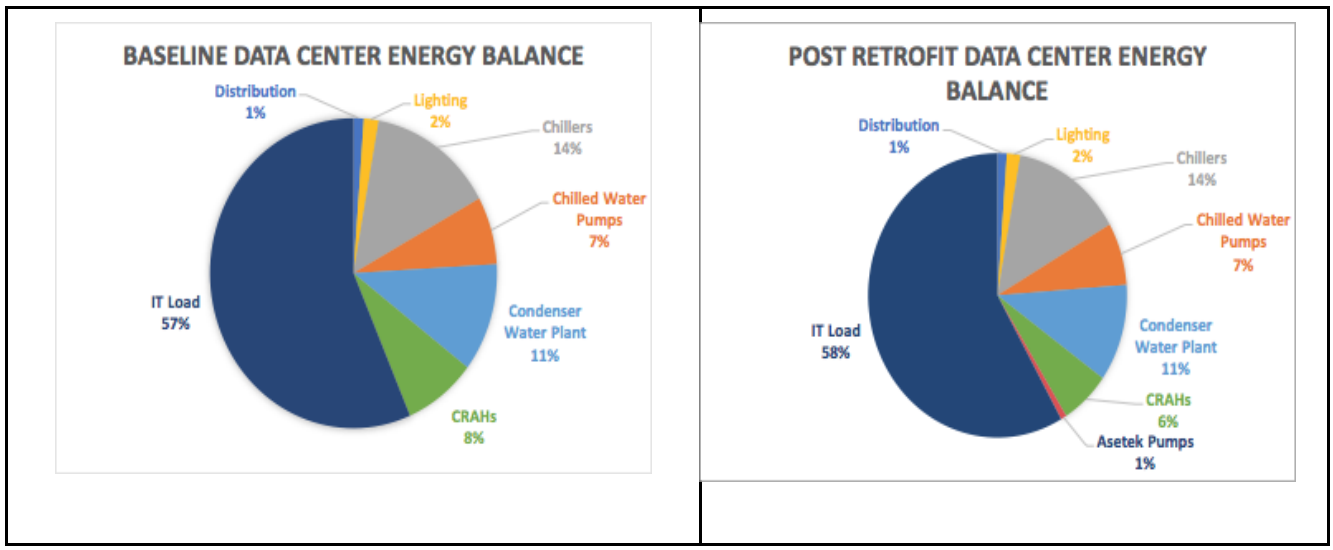


Figure 2: (a) Baseline Full Data Center Energy Balance (kWh) (left); (b) Post Retrofit Data Center Energy Balance (kWh)

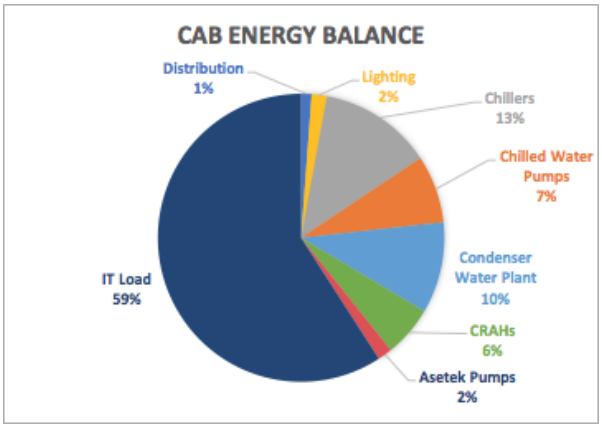


Figure 3: Post Retrofit Cab (kWh)

Table 2 below summarizes the PUE and pPUEs defined above for the proposed upgrade.

Table 2: PUE and pPUE Summary

| Baseline | Post Retrofit | | | |
|---|---|---|---|---|
| Data Center PUE | Data Center PUE | Cab pPUE | Asetek w/ CHW pPUE | Asetek Only pPUE |
| 1.70 | 1.66 | 1.51 | 1.44 | 1.01 |

The total overall data center PUE reduced from 1.70 to 1.66 as a result of the retrofit. The "Asetek Only" pPUE that only includes the CDU circulating pump and Asetek rack pumps are very efficient at removing heat from the servers they serve is calculated to be 1.01. However, since the Asetek system rejects heat to the building chilled water system, once the chilled water and condenser water plant efficiencies are considered, the pPUE is much higher and this "Asetek w/CHW" pPUE was calculated to be 1.44. Based on the data we have available, the combined chilled water and condenser water plant efficiency is relatively low. Rejecting heat directly to an efficient condenser water loop would reduce the PUE significantly.

The Cab pPUE (1.51) isolates Cab and its associated cooling systems, including both CRAH fans and the Asetek system. It is our best estimate for what the PUE of the data center would be if all servers were cooled by Asetek. We expect this reduction to have been much more significant if the Asetek system was not required to reject heat to the chilled water system.

**Top-down Statistical Model Development**

Based on engineering principles, the power required for the mechanical system to cool the data center is a function of the cooling load and the outside weather conditions. The internal load is dominated by the IT load from various clusters including Cab. The IT power is also gathered for other clusters that are air-cooled. Based on the regression analysis, the model (Equation 1) below is developed to predict the mechanical power (Y) needed for the computer racks for a given weather condition.

$$\mathbf{Y = 43.1218 + 0.1525*(IT\ Load\ 1) + 0.2915*(IT\ Load\ 2) + 0.5991*OAT} \quad \text{----------(1)}$$

The regression coefficients shown in Equation 1 represent the change in the dependent variable resulting from a one-unit change in the predictor variable, all other variables being held constant. In the regression model, for example, a unit increase in IT Load 1, that includes the power draw by the Cab cluster- increases mechanical power to cool the cluster by about 15%, while a unit increase in IT Load 2- the air-cooled IT equipment- increases the mechanical power by nearly 30%.

Table 3 indicates the summary to assess goodness of fit for the developed model. One of the key metrics is Coefficient of Determination, $R^2$, which measures the extent to which variations in the dependent variable y can be explained by the regression model. Root mean squared error (RMSE) or Standard error of the estimate (SE) is an indicator of the scatter, or random variability, in the data, and hence is an average of how much an actual Y-value differs from the predicted Y-value. It is the standard deviation of errors of prediction about the regression line. Coefficient of variation of the root mean squared error CV(RMSE) is the RMSE normalized by the average y-value. Net determination bias error (NBE) is simply the percentage error in the energy use predicted by the model compared to the actual energy use.

Table 3: Summary of the model fit characteristics

| Statistic | Description | Model |
|---|---|---|
| Root mean squared error (RMSE) | An indicator of the scatter, or random variability, in the data, and hence is an average of how much an actual y-value differs from the predicted y-value | 11.45 |
| CVRMSE | Non-dimensional Metric that normalizes RMSE by the average y-value that describes how well the model fits the data. | 4.4% |
| Net Determination Bias Error (NBE) | Percentage error in the energy use predicted by the model compared to the actual energy use. The sum of the differences between actual and predicted energy use should be zero | 0.01% |
| Coefficient of determination ($R^2$) | Measures the extent to which variations in the dependent variable y can be explained by the regression model | 74.3% |

**Savings Analysis**

The developed baseline model was used to project the load consumed by the mechanical system that includes both the chiller and the CRAH systems for the post retrofit conditions. The data from 2017 was used to represent the post retrofit period and the conditions that included the IT power and outside air temperature were used to project baseline load for the overall mechanical system for each of the 30 minute intervals. This projected baseline mechanical load (blue in Figure 4) is assumed to be the load as if no retrofit was implemented. This predicted mechanical baseline load is compared with corresponding actual mechanical load from 2017 (red in Figure 5) to assess the impact of the Asetek retrofit.

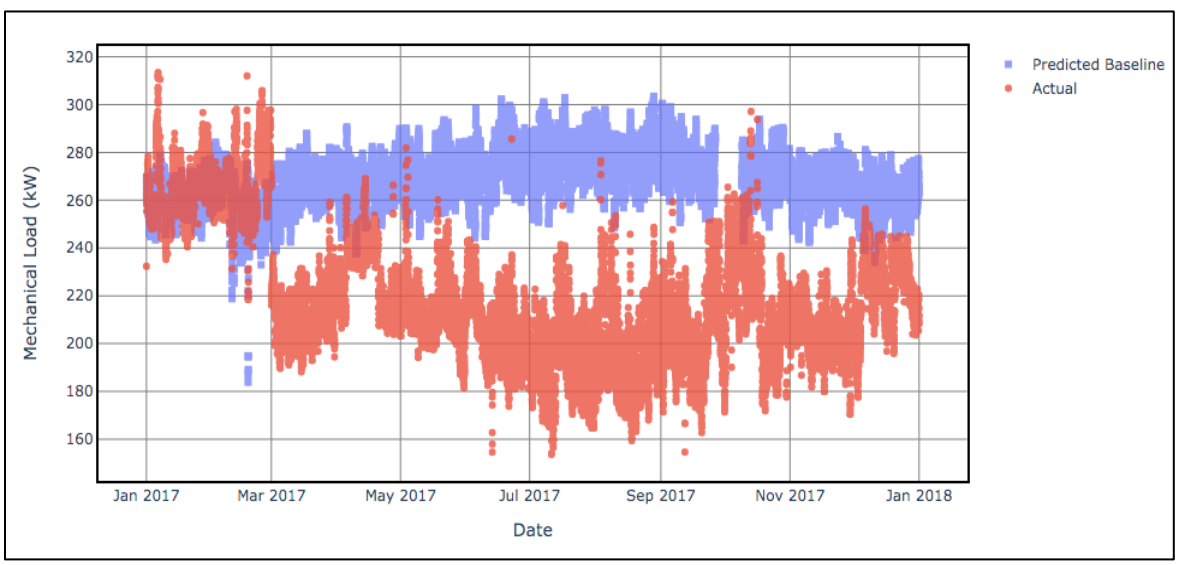


Figure 4: Plot showing the predicted baseline and actual mechanical power for the post retrofit period

**Results**

The actual mechanical load for the post retrofit conditions is unusually high for the first two months in 2017 (Figure 4) and no plausible explanation was obtained from any of the implementation team or the site team members. These data were retained and not excluded from the analysis. This predicted baseline mechanical load and the actual mechanical load, including

the other mechanical and office loads, were analyzed to calculate the PUE for the baseline and post retrofit IT and weather conditions. The PUE is calculated as:

$PUE_{BL}$ = [Total IT Load + Total Predicted Baseline Mechanical Load + Total Office Load+ Total Condenser Predicted Water Load]/[Total IT Load]

$PUE_{PR}$ = [Total IT Load + Total Actual Mechanical Load + Total Office Load+ Total Condenser Actual Water Load]/[Total IT Load]

Total IT Load (kW) = IT Load 1 (kW) + IT Load 2 (kW) which is assumed to be the post retrofit IT load.

Total Office Load is the load from office rooms; lighting and plug loads, assumed to not change between baseline and post retrofit conditions.

Total Condenser Predicted Water Load (kW) = [Total IT Load (kW) + Total Predicted Baseline Mechanical Load (kW) + Total Office Load (kW)]*3,412 [Btus/kWh]/12000 [Btu/ton-hr)] * Condenser Water kW/ton[1]

Total Condenser Actual Water Load (kW) = [Total IT Load (kW) + Total Actual Mechanical Load (kW) + Total Office Load (kW)]*3,412 [Btus/kWh]/12000 [Btu/ton-hr] * Condenser Water kW/ton[1]

The baseline PUE ($PUE_{BL}$ - blue line in Figure 5) is compared with post-retrofit PUE ($PUE_{PR}$) conditions (red line in Figure 6). The average annual PUE for the baseline conditions dropped from 1.60 to 1.53 for the post retrofit conditions, which is assumed attributable to the Asetek implementation. Also, none of the energy consumption related to Asetek equipment is included in this analysis, which when included would increase the post-retrofit PUE slightly.

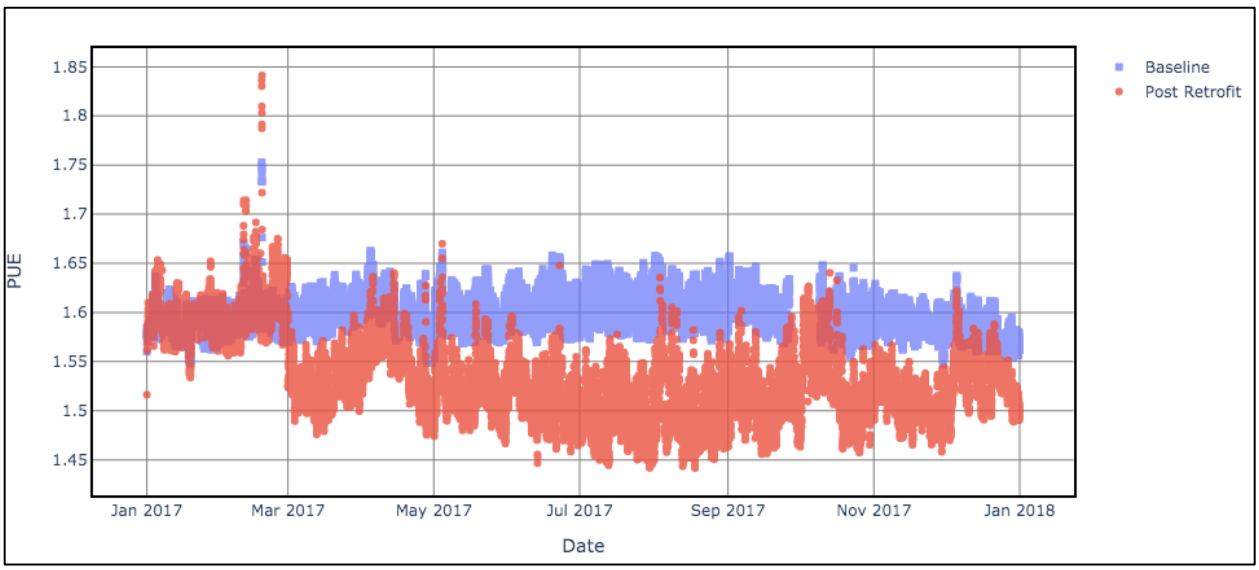


Figure 5: Plot showing the predicted baseline and actual PUE for the post retrofit period.

[1,2] Estimated by site as 0.58 and assumed to remain constant and not affected by Asetek

The heat gathered by all the Asetek RackCDUs was compared with the Cab power for a 10-day period in February of 2018. This heat gathered by CDUs was found to be 37% of the overall Cab IT power, which is considerably lower than what was found in previous research. For example, a study done in 2014 by LBNL (Coles, 2014) found heat capture rates by the Asetek system ranged from about 47% to 63% (depending on IT load) at the same cooling water supply temperature to the rack CDUs as in the LLNL case. However, direct comparisons are complicated by the differences between the LBNL study and LLNL shown in Table 4.

Table 4: Comparison of conditions at LLNL and LBNL 2014 study

| Parameter | LBNL 2014 | LLNL |
|---|---|---|
| **Water-cooled components** | Processors and memory | Processors only |
| **Supply water temperature** | 15, 20, 25, 30, 40, and 45° C | 26° C average |
| **Water flow rate per rack** | 4.9 gpm | 3.9 gpm average |
| **Server power** | 120, 270, and 430W each | 270 W average |
| **Room air temperature** | 28° C | <22° C average |

The 2014 configuration included water-cooled memory and used higher water flow rates that at LLNL, both of which would result in higher capture rates at LBNL. The room temperature was higher in 2014 than at LLNL (the 22° C is the average of the CRAH return temperatures, with the CRAH supply temperatures averaging about 16° C), also resulting in higher capture to water in 2014. The HPC IT equipment and software used for loading the IT equipment also was different, so the fraction of heat going to processors and memory would differ, but not enough is known to speculate on which direction these differences would affect the capture rate. In another study (Sickinger, 2014), the National Renewable Energy Laboratory found 48% of the IT load rejected to water with a system using water-cooled CPUs (only, not memory chips), 1.1 gpm of 17° C cooling water supplied to the RackCDUs, and 20° C inlet air temperature to the IT equipment. The lower water flow was used to make the return water warm enough for heat reclaim, making this condition comparable to the approximately 58% capture in the LBNL test under very similar water temperatures, load, and flow, but with the memory chips also water cooled.

## Summary and Conclusions

The top-down analysis that includes high-level meter data for IT loads, mechanical loads including both the CRAHs and chillers, and other ancillary loads indicates that the overall PUE dropped from 1.60 for an all air-cooled system to 1.53 for a system where a portion of the servers are equipped with a water–based Asetek cooling system. From the bottom-up analysis, using a bin-data analysis to understand the performance of the Asetek system, the total overall data center, where Cab is part of the load, the overall PUE reduced from 1.70 to 1.66. However, when analyzing the performance of Asetek cooling system that only includes the CDU circulating pump, and air-cooling systems for Cab, the pPUE is estimated to be 1.51 which would be the expected PUE of the data center if all the servers were cooled by Asetek. We expect this

reduction to have been significantly larger if the Asetek system was not required to reject heat to the chilled water system; see further discussion below. The overall heat rejection fraction from Cab to the Asetek system was 37%, lower than in previous studies.

**Further Observations**:
It should be noted that the LLNL installation was not a good example of the potential of this water-cooling strategy. At LLNL, the heat was rejected to the same chilled water system whether air or water cooled. Because the Asetek system provides much better thermal coupling from the chip to the heat rejection system, substantially warmer water could be used, for example from the condensing water system. Doing so would eliminate the chiller and chilled water pumps from the heat rejected by the Asetek system, significantly improving efficiency and dropping PUE. LBNL's study (Coles, 2014) found, at 50% server load, 14% overall site energy reduction for dry cooler rejection with chiller boost (with 40°C supply water—minimizing water use for cooling and maximizing the potential for heat recovery) and 20% overall site energy reduction for cooling tower heat rejection with chiller boost (with 20°C supply water, minimizing the overall energy use). The cooling pPUEs were 1.34 for the base case, 1.24 for the dry cooler case, and 1.20 for the cooling tower case. Other configurations and loads were modeled in the study. The point is that by taking full advantage of the Asetek water cooling system, one can greatly reduce or eliminate the use of the chiller plant, significantly increasing the savings relative to that of the LLNL case study.

In the LLNL system, an intermediate heat exchanger (the CDU) was used between the primary cooling water (i.e. the chilled water) and the cold side of the RackCDU (the hot side of which is where the heat is gathered from the individual IT cooling circuits). Such heat exchangers provide isolation between the water loops, the primary advantage of which is minimizing water release in the event of a leak, and can provide better control stability. However, such exchangers create the need for an additional pump (since an additional water loop is created) and additional pumping energy requirements to overcome the pressure drops on both primary and secondary sides of the heat exchanger, and they impose an additional temperature difference, requiring the cooling water to be typically 2°F colder for the same temperature at the IT load, decreasing the efficiency of the heat rejection when compressor cooled and reducing the number of compressor-free hours (or forcing less-efficient operation) in systems with water-side economizers. Thus, in general, eliminating the intermediate exchanger would save capital and operating costs, the latter in the form of energy savings.

In this case, the standard CDU configuration was used for design simplicity warranted by a temporary installation. Also because the chilled water temperature was not allowed to vary, and because no compressor-free cooling option was available, very little energy savings would have resulted. Other means to mitigate potential leaks (like a leak detection system with alarms and control valves that would automatically close to isolate the leak) could be provided in systems without an intermediate heat exchanger.

In short, the Asetek liquid cooling system and associated facility heat-rejection systems resulted in some energy savings at LLNL (3-5% overall for the data center, and 7% for the Cab cluster). In more typical applications, where advantage can be taken of the Asetek system to reduce or eliminate chiller operation, much higher savings can be anticipated.

**Acknowledgements**

This work was performed by the Lawrence Berkeley National Laboratory and kW Engineering under contract to Asetek (San Jose, California) and supported by the California Energy Commission's Public Interest Energy Research Program and by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy under Contract No. DE---AC02---05CH11231. Thanks to key partner Lawrence Livermore National Laboratory for the High-Performance Computing clusters used in the study, for the facilities that housed and cooled them, and for the metering and data acquisition used for gathering the data. Specific appreciation goes to Steve Branton and Sean Ferguson of Asetek and Anna Maria Bailey, Ghaleb Abdulla, John Weaver, and Trent D'Hooge of LLNL.

**Disclaimer**

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California. Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.

## References

Capozzoli A., and G. Primiceri. 2015. "Cooling Systems in Data Centers: State of Art and Emerging Technologies". Elsevier. Sustainability in Energy and Buildings: Proceedings of the 7th International Conference SEB-15. 83, p.484–93.

Coles, H., and S. Greenberg. 2014. "Direct Liquid Cooling for Electronic Equipment". Berkeley CA. Lawrence Berkeley National Laboratory Report 6641E.

Patankar, S. V. 2010. "Airflow and Cooling in a Data Center". New York. American Society of Mechanical Engineers. ASME J. Heat Transfer, 132(7), p. 073001-17.

Shehabi, W.L.A., S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner. 2016. "United States Data Center Energy Usage Report". Berkeley CA. Lawrence Berkeley National Laboratory Report 1005775.

Sickinger, D., O. Van Geet, and C. Ravenscroft. 2014. "Energy Performance Testing of Asetek's RackCDU System at NREL's High Performance Computing Data Center", Golden, CO. National Renewable Energy Laboratory, NREL/TP-2C00-62905.