CALIFORNIA ENERGY COMMISSION

CALIFORNIA NATURAL RESOURCES AGENCY

ENERGY RESEARCH AND DEVELOPMENT DIVISION

FINAL PROJECT REPORT

Demonstration of Low-Cost Data Center Liquid Cooling

June 2024 | CEC-500-2024-061

PREPARED BY:

Steve Branton
Shankar Earni
Sean Ferguson
Asetek Inc
Steve Greenberg
Brian Park
Primary Authors

Kevin Mori
Project Manager
California Energy Commission

Agreement Number: EPC-14-088

ACKNOWLEDGEMENTS

This work was performed by the Lawrence Berkeley National Laboratory and kW Engineering under contract to Asetek (San Jose, Santa Clara County, California) and supported by the California Energy Commission's Electric Program Investment Charge 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 our 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 Lawrence Livermore National Laboratory.

PREFACE

The California Energy Commission's (CEC) Energy Research and Development Division supports energy research and development programs to spur innovation in energy efficiency, renewable energy and advanced clean generation, energy-related environmental protection, energy transmission, and distribution and transportation.

In 2012, the Electric Program Investment Charge (EPIC) was established by the California Public Utilities Commission to fund public investments in research to create and advance new energy solutions, foster regional innovation, and bring ideas from the lab to the marketplace. The EPIC Program is funded by California utility customers under the auspices of the California Public Utilities Commission. The CEC and the state's three largest investor-owned utilities — Pacific Gas and Electric Company, San Diego Gas & Electric Company, and Southern California Edison Company — were selected to administer the EPIC funds and advance novel technologies, tools, and strategies that provide benefits to their electric ratepayers.

The CEC is committed to ensuring public participation in its research and development programs that promote greater reliability, lower costs, and increased safety for the California electric ratepayer and include:

- Providing societal benefits.
- Reducing greenhouse gas emissions in the electricity sector at the lowest possible cost.
- Supporting California's loading order to meet energy needs, first with energy efficiency and demand response, next with renewable energy (distributed generation and utility scale), and finally with clean, conventional electricity supply.
- Supporting low-emission vehicles and transportation.
- Providing economic development.
- Using ratepayer funds efficiently.

For more information about the Energy Research and Development Division, please visit the CEC's research website (www.energy.ca.gov/research/) or contact the Energy Research and Development Division at ERDD@energy.ca.gov.

ABSTRACT

Data centers consume approximately 2 percent of California's electricity demand. Reports by Digital Power Group at the time of this application indicate that the amount could be as high as 10 percent. While overall commercial and residential building energy efficiency has made dramatic advances over the past 5 to 10 years, efficiency in data centers continues to increase in size and power, and the percentage of California electricity consumed by this single industry is increasing. Approximately 40 percent of this electricity is used to cool the data centers. As such, data center cooling efficiency represents one of the largest and most important energy-efficiency measures in the state.

RackCDU™ is a unique, pre-commercial data center efficiency technology that brings high-performance liquid cooling directly to the hottest elements inside each server, with the potential to cut cooling energy by 60 percent to 80 percent and server energy consumption by an additional 5 percent to 10 percent. This innovative design could be retrofitted into existing servers and data centers, enabling rapid adoption across all California data centers. If deployed in existing data centers, RackCDU™ could annually save California ratepayers up to 2,400 gigawatt hours of electricity and $340 million.

This project validated the performance, reliability, and lifecycle cost benefits of RackCDU™ in two full-scale California data centers. This project provided operational insights to create awareness and acceptance of this advanced pre-commercial technology across California, accelerating commercialization and leading to significant ratepayer benefits.

Keywords: liquid cooling, RackCDU™, data centers, servers

Please use the following citation for this report:

Branton, Steve, Shankar Earni, and Steve Greenberg, 2020. *Demonstration of Low-Cost Data Center Liquid Cooling*. California Energy Commission. Publication Number: CEC-500-2024-061.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Executive Summary

Introduction

California is home to many data centers, which consume a disproportionate amount of electricity. Approximately 40 percent of the electricity used in data centers is used for cooling; and, as traditional building efficiency improves, data centers continue to grow in size and power. Improving data center cooling efficiency represents one of the major energy-efficiency measures for this sector. The standard way of cooling data center racks is with high amounts of air flow, and these require large amounts of electricity to run air conditioning for the computing spaces. Liquid cooling has been used for decades at the computer level and could be scaled for data center applications to reduce electricity consumption.

Project Purpose

The goal of this project is to validate the performance, reliability, and lifecycle cost benefits of RackCDU™ in two full-scale California data centers. RackCDU™ is a component of the Asetek system. It is a unique, pre-commercial data center efficiency technology that could reduce cooling energy by 60 percent to 80 percent and server energy consumption by an additional 5 percent to 10 percent. The project included 18 months of pilot-scale testing to generate the full-scale data and operational insights needed to create awareness and acceptance of this advanced pre-commercial technology across California, accelerate commercialization, and document ratepayer benefits.

The research team's project objectives included the following:

1. Install RackCDUs™ in a full-scale retrofit at two California data centers without disrupting operations.
2. Study and optimize integrated system performance under different operational conditions.
3. Demonstrate and document, through independent measurement and verification, 60 percent to 80 percent cooling energy savings, 5 percent to 10 percent server energy savings, and 20 percent to 30 percent peak-load savings.
4. Quantify non-energy benefits, including improved computation performance and lower noise.
5. Demonstrate total system reliability at or above pre-retrofit levels.
6. Demonstrate projected lifecycle energy and greenhouse gas savings of 20 percent to 30 percent, payback of less than 12 months, and a lifetime return on investment of more than six times at each site.
7. Demonstrate operational transparency by users, sufficient to overcome any perceived barriers to broad-scale adoption of this technology.

8. Communicate results of this project broadly to California data center owners and policy makers.

Project Approach

To achieve the stated objectives, two clusters of high-performance computing equipment were equipped with Asetek's water cooling systems. These included the Cabernet (Cab) and the Topaz high-performance computing clusters. Both were retrofitted with the Asetek system. The Asetek system is a hybrid scheme with water-cooled cold plates on high-heat-producing components, with the remainder of the heat being removed by conventional air-cooling systems. Installations were performed through a combination of "rolling retrofits" of existing servers and pre-installed cooling loops into new servers prior to server installation at the data centers. This included retrofitting RackCDU™ internal loops into existing operational servers a few at a time, while using load-shifting procedures to keep the rest of the data center in full operation.

To determine energy savings of the Asetek system, data were gathered and analyzed using statistical and engineering models, as described in more detail in Options A through D in Chapter 2, "Project Approach — Measurement and Verification Plan." The two high-performance computing clusters were in two different buildings (Building 451 and Building 654) at the Lawrence Livermore National Laboratory in California.

In addition to energy performance parameters, the Asetek team also monitored and estimated key server performance parameters (central processing unit leakage and power capping and soft-data error rates) and total server and system reliability, including server up-time before and after RackCDU™ installation. The data generated were used to calculate lifetime energy savings and the total cost of ownership at each site. This was then used to calculate lifetime greenhouse gas emission effects of RackCDU™ and to extrapolate the potential effect of RackCDU™ on energy, cost, and greenhouse gas emissions across the full California data center stock.

Project Results

The statistical analysis in Building 451 showed 5 percent overall energy savings for the data center, and power usage effectiveness, which is the ratio of total data center energy to information technology energy, dropped from 1.60 to 1.53 (lower is better). The engineering analysis showed 3 percent overall energy savings (power usage effectiveness from 1.64 to 1.59). This translates to a carbon dioxide emission reduction of 0.3 metric tons per kilowatt of information technology load retrofitted, using the California average carbon dioxide emissions per kilowatt-hour generated. Greater savings, on the order of 10 percent to 15 percent, would be possible if the chilled water system was not used for rejecting the heat from the Asetek system. About 37 percent of the heat from the Cab system (a high-performance computing cluster) was rejected to the cooling water, lower than at other installations. The primary benefit of liquid cooling in this project was the reduction in overall temperature and improved temperature uniformity across the rack.

The air-cooling system in Building 654 was very efficient, more efficient than the water-cooling system in that building. This unexpected finding was due to turn-down issues and the control scheme for the water-based system, as the facility's water system was substantially oversized for the load provided by the test cluster (Topaz). The partial power use effectiveness of the air-cooling system for the control cluster (Quartz) was 1.07; the partial power use effectiveness of the hybrid cooling system for Topaz was 1.32, for an overall facility power usage effectiveness of 1.13. The partial power usage effectiveness defines a certain portion of the overall power use effectiveness of a data center within a clearly defined boundary — in this case, the boundaries were for the liquid-cooled Topaz cluster and the air-cooled Quartz cluster. Thus, the overall energy use intensity was 23 percent higher for the water-cooled case than for the air-cooled system. The heat capture fraction of 23 percent was also considered low compared to other facilities. If the conditions noted are addressed, the water-cooling system is expected to outperform the air-cooling system. However, it is unlikely that savings greater than 50 percent could be achieved, and, given the already excellent cooling efficiency in this building, the total power savings would be small.

Technology/Knowledge Transfer/Market Adoption (Advancing the Research to Market)

Knowledge gained from the project has been made available to the public through a variety of sources, including the Asetek website, conferences, presentations, and discussions with partners, utilities, and customers. Asetek has discussed the lessons learned with data center operators around the state and the country in person, with the goal of implementing lessons learned in future installations.

Benefits to California

Implementation of the proposed technology will benefit investor-owned utility electricity ratepayers in the near term by reducing the electricity bill for their data centers by up to 10 percent. Greater savings of 20 percent to 30 percent may be possible if the chilled water system was not used for rejecting the heat from the Asetek system. Although the researchers did not perform peak load reduction analysis in this project, RackCDU™ could reduce grid strain during peak times and improve grid reliability and safety, by cost effectively reducing statewide data center energy consumption. This technology has the potential to cut statewide commercial electricity consumption by more than 2 percent, saving California 2,400 gigawatt hours of electricity and $340 million each year.

CHAPTER 1: Introduction

The initial value proposition that drove this study was that data centers consume 2.2 percent of all electricity nationally, with a disproportionate amount of this energy coming from California (recent reports by Digital Power Group indicate that it could be as high as 10 percent). While overall commercial and residential building energy efficiency has made dramatic advances over the past 5 to 10 years, data centers continue to grow in size and power, and the percentage of California electricity consumed by this single industry is increasing; approximately 40 percent is used to cool these data centers.

As such, data center cooling efficiency represents one of the largest and most important energy-efficiency measures in the state. RackCDU™ is a unique, pre-commercial data center efficiency technology that brings high-performance liquid cooling directly to the hottest elements inside each server ("hot-spot cooling"). The technology has the potential to cut cooling energy by 60 percent to 80 percent and cut server energy consumption by an additional 5 percent to 10 percent. This innovative design could be retrofitted into existing servers and data centers, enabling rapid adoption across all California data centers

In this project, the Asetek team demonstrated RackCDU™ in two California data centers, at full scale, to create awareness and generate all of the performance, reliability, and usability data needed to catalyze adoption across the state. The goal was to validate the performance, reliability, and lifecycle cost benefits of RackCDU™ in two full-scale California data centers. The project included 18 months of pilot-scale testing to generate the full-scale data and operational insights needed to create awareness and acceptance of this advanced pre-commercial technology across California.

Project objectives included the following:

1. Install RackCDUs™ in a full-scale retrofit at two California data centers without disrupting operations.
2. Study and optimize integrated system performance under different operational conditions.
3. Demonstrate and document, through independent measurement and verification (M&V), 60 percent to 80 percent cooling energy savings, 5 percent to 10 percent server energy savings, and 20 percent to 30 percent peak-load savings.
4. Quantify non-energy benefits, including improved computation performance and lower noise.
5. Demonstrate total system reliability at or above pre-retrofit levels.
6. Demonstrate projected lifecycle energy and greenhouse gas (GHG) savings of 20 percent to 30 percent, payback of less than 12 months, and a lifetime return on investment (ROI) greater than six times at each site.

7. Demonstrate operational transparency by users, sufficient to overcome any perceived barriers to broad-scale adoption of this technology.
8. Communicate results of this project broadly to California data center owners and policy makers.

Facilities Descriptions

This study was conducted in two different data centers located in two separate buildings on the campus of Lawrence Livermore National Laboratories (LLNL). The existing air-cooling systems in these two buildings are radically different, one (Building 451) representing a typical legacy data center cooling system and one (Building 654) representing a data center with a state-of-the-art air-cooled system.

Building 451 Description

The data center in Building 451 (B451) is supplied with conditioned air through a raised floor with a 3-foot underfloor plenum (an air-filled space in a structure that receives air from a blower for distribution) and with open return.

Computer Room Air-Handling Units

The data center is conditioned primarily by 37 Data Aire computer room air-handling (CRAH) units with chilled water (CHW) system coils and fans equipped with variable frequency drives (VFDs). The CHW coils in the CRAH units have three-way valves, controlled to a constant-supply air temperature by a building automation system made by Automated Logic Corporation. The fan VFDs are controlled based on return air temperature.

Chilled Water System

There are three Trane, centrifugal, water-cooled chillers that provide CHW to the CRAH units as well as to the cooling distribution unit, supplying closed-loop cooling water to the Asetek RackCDU™ (described in Table 1). Trend data show a constant CHW supply temperature (CHWST) setpoint of 46 degrees Fahrenheit (°F [8 degrees Celsius [°C]). The CHW system has a primary-only pumping configuration, with four constant speed CHW pumps (Table 2). During Asetek's site visit, two of the four pumps were running.

Table 1: Chiller Summary

Manufacturer	Model	Qty	Type of Compressor	Year Installed
Trane	CVHF-077N-AW003	3	Centrifugal	1997

Source: Asetek, Inc.

Table 2: Chilled Water Pump Summary

Manufacturer	Series	Qty	Rated Motor Horse-power (hp)	Flow Rate (gpm)	Head (ft)	Pump Efficiency	Motor Efficiency
Bell & Gossett	VSC 9.875 BF LHR	4	50	1535	80	80 percent	94.5 percent

Gpm = gallons per minute
Source: Asetek, Inc.

The CHW system serving the data center also serves the following loads:

- Cabernet (Cab) high-performance computing cluster retrofitted with Asetek system
- Two cooling distribution units (CDUs) feeding other computer clusters
- Water-cooled testbeds for clusters (via heat exchanger)

The chillers are served by the campus condenser water loop. There are no condenser water pumps in the building.

Humidity Control

Humidification and dehumidification in the data center are achieved through two make-up air units (Table 3) that supply a total of 15,000 cubic feet per minute (cfm) of outside air. Relative humidity is kept within roughly 35 percent to 45 percent. Since the CHWST is 46°F (8°C), there is also some natural dehumidification at the CRAH.

Table 3: Make-Up Air Unit Summary

Manufacturer	Supply Fan Motor Size (hp)	Rated Airflow (cfm)	Total Static Pressure (IWC)	Fan Brake Horsepower
Pace	10	8,700	3.25	6.82

Source: Asetek, Inc.

Uninterruptible Power Supplies

Uninterruptible power supplies serve only a small amount of information technology (IT) equipment.

Energy Conservation Measure Description

The data center has one high-performance computing cluster, Cab, which was retrofitted with a liquid-cooled Asetek system for this demonstration (see metering diagram at Figure 1). The Asetek system uses direct liquid cooling to cool some components of computer servers. Direct cooling brings water to high heat, producing electronic components (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 baseline condition uses conventional air cooling. The liquid cooling system uses two 4-watt (W) pumps for each server to circulate a glycol and water solution directly through the server. There are 2464 of these pumps serving Cab. A CDU at each rack (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 CHW system. By transferring heat from the servers to the CHW loop this way, the heat bypasses the CRAH units, saving CRAH fan energy, resulting in reductions in chiller energy. There are no energy savings from the CHW pumps because the pumps are at constant speed and the CRAH control valves are three-way valves, resulting in constant flow and pump energy regardless of load.

Figure 1: Existing Building 451 Level Power Metering

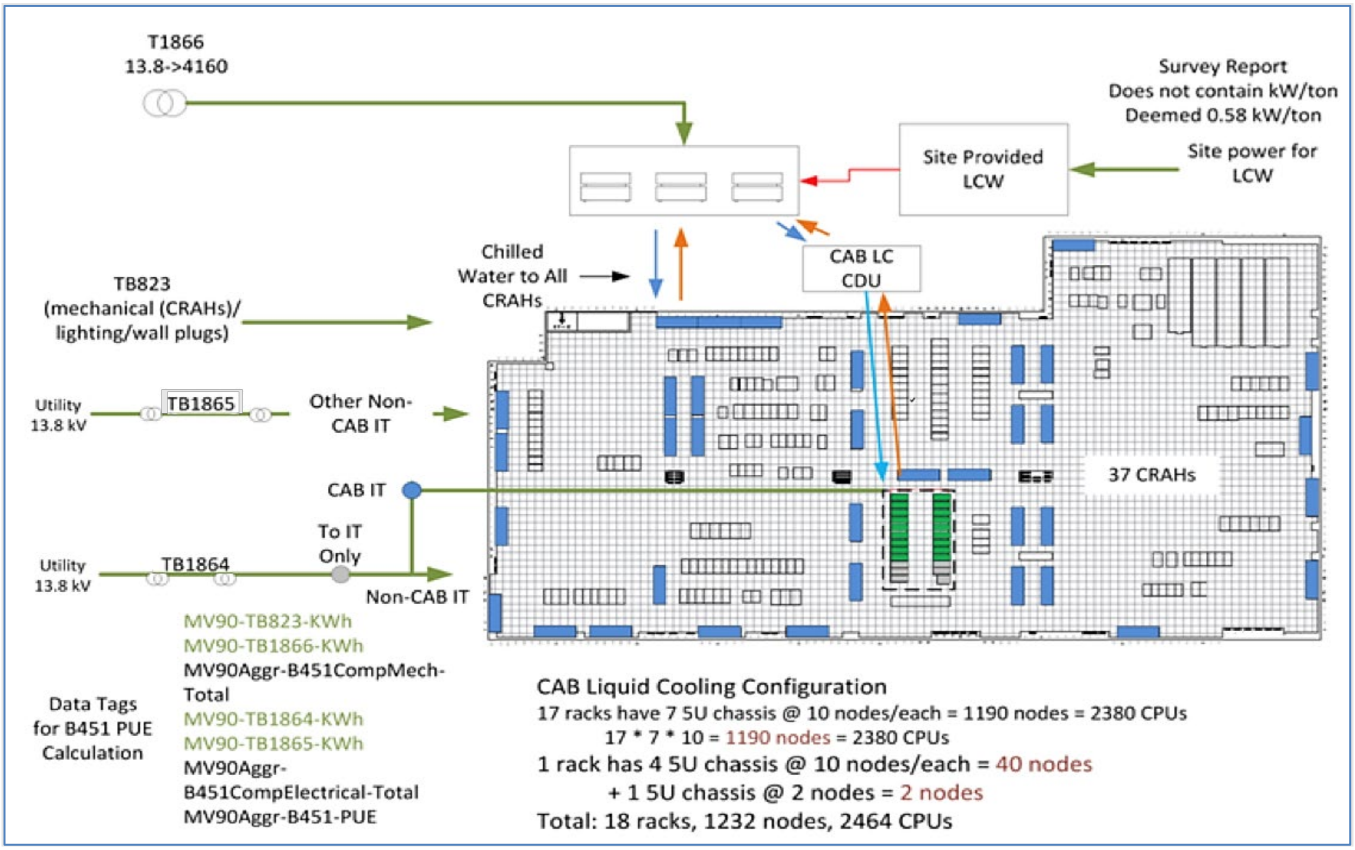

Source: Asetek, Inc.

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. This 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 CHW system directly, compared to relying completely on the existing cooling system that uses only air and CRAH. 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 CHW system.

Building 654 Description

Building 654 housed two high-performance computing clusters at the time of the study, one served by an Asetek system (Topaz), and one served by an underfloor air-cooling system (Quartz).

Topaz is located on the first floor and is served by the liquid-cooled Asetek system and its associated central CDU, pumps, and outdoor fluid cooler for heat rejection from the water. The first floor serves as an air supply plenum for the second floor, where the air is used by an air-cooled system (Quartz), as shown in Figure 2.

The second floor has partial hot-aisle containment. The return air is not ducted directly from the hot aisles, so it draws air from both the hot and the cold aisles. This suboptimal air management situation makes it impossible to control pressure appropriately, resulting in a relatively high airflow requirement for the IT load served.

Humidity Control

The building has no humidity control.

Air-Handling Units

The data center is conditioned by two air handling units (AHUs), with 100 percent direct evaporative cooling capability (no compressor-based CHW or direct expansion cooling). The units are equipped with direct expansion coils that could be used as a backup for cooling, but there are currently no compressors attached to them.

The AHUs have two totally separated sides: a waterside and an airside.

- The waterside cools process water used by the Asetek system by blowing evaporatively cooled outside air through a process water coil. After it passes the coil, all the air is exhausted out at the back of the AHU. The waterside section is essentially an air-cooled fluid cooler with a direct evaporative pre-cooler on the cooling air. This system is not configured to allow natural draft cooling when air temperatures are low.
- The airside provides evaporatively cooled outside air directly to the first floor. Conditioned air from the AHUs is supplied to the first floor and flows up to the second floor through perforated tiles. Return air from the second floor is ducted back to the AHUs for some combination of recirculation and exhaust at the units, depending on outside air temperature (OAT).

Figure 2: Schematic of Section of Building 654 Data Center

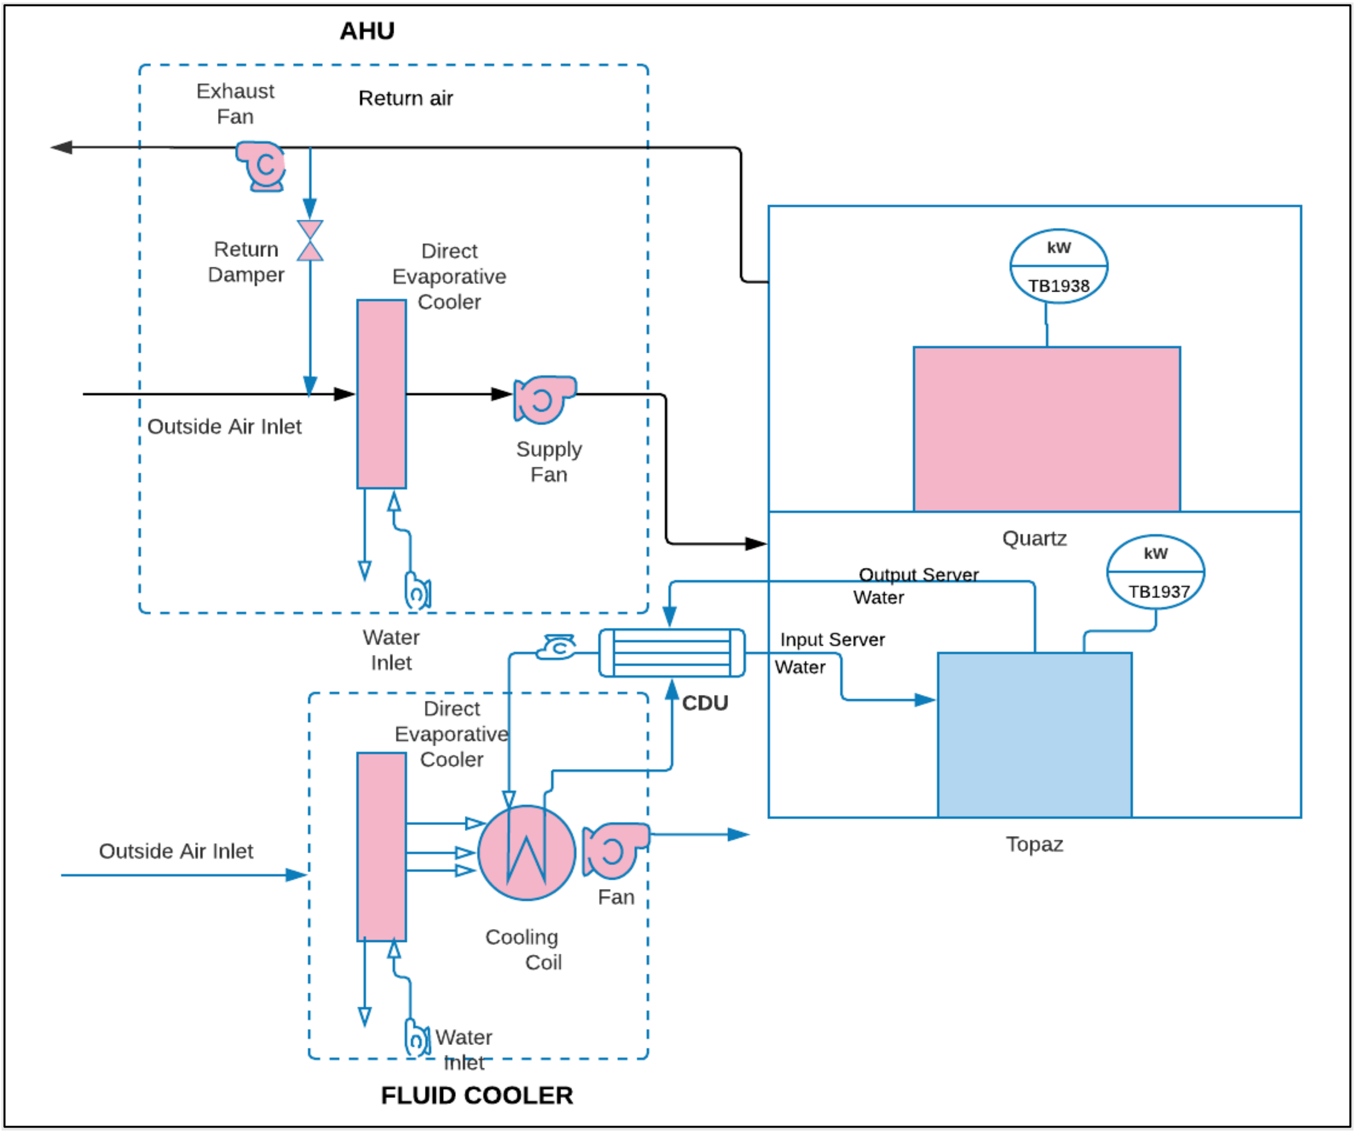

Source: Asetek, Inc.

Table 4: Air Handling Unit Summary

Manufacturer	Supply Fan Motor Size (hp)	Fluid Cooler Fan Motor Size (hp)	Exhaust Fan Motor Size (hp)	Rated Airflow (cfm)
Basics	100	100	(3) 10	80,000

Source: Asetek, Inc.

Each AHU has one 100-hp supply fan and three 10-hp exhaust fans on the airside, and one 100-hp fluid cooler fan on the waterside (Table 4). All fans have VFDs. On the airside, the supply fans are set to maintain a constant flow rate. The exhaust fans are controlled to maintain a return duct static pressure. In doing so, they modulate to match the quantity of outside air introduced by the supply fans. On the waterside, the fluid cooler fan modulates to maintain water temperature in the primary cooling water loop.

According to the facility's staff, the exhaust fans are too small to exhaust all the air that the supply fans can provide. Therefore, the supply fans must run at lower speeds to balance the airflow. The site is planning to retrofit the exhaust fans with bigger motors to increase flow for both the exhaust and the supply sides.

The modeling work was done by Lawrence Berkeley National Laboratory (LBNL) in conjunction with kW Engineering. During a monitoring and verification site visit, the modeling team observed a supply air temperature of 57°F (14°C). According to the facilities staff, it is challenging to condition the air on days with high wet-bulb temperatures, which means the evaporative cooling is limited as to how much cooling can be provided. When the outside wet-bulb temperature is too high, the computers are throttled (and are sometimes shut down) to protect them. The highest wet-bulb temperature the staff had experienced by the time of the study was 74°F (23°C). The site had managed to provide 78°F (26°C) air to the space on a 112°F (44°C) (dry-bulb temperature) day during the summer of 2018.

Process Water Loop

Process water used in the central CDU feeding the Asetek RackCDUs™ is circulated from the waterside of the AHUs via two process loop pumps with VFDs. During a site visit, only one pump was running, at 45 hertz (Table 5).

Table 5: Process Water Pump Summary

Manufacturer	Series	Qty	Rated Motor Horsepower (hp)	Flow Rate (gpm)	Head (ft)	Motor Efficiency
Bell & Gossett	e-1510 SSF 11.875	2	40	800	120	94.1 percent

Source: Asetek, Inc.

Asetek Liquid-Cooled System

The system installed in the first floor Topaz cluster at Building 654 has one central CDU that provides cooling to the racks and isolates the racks from the main facilities loop. There are 1,488 4-W server-level pumps serving Topaz. There is one constant-speed pump circulating water from the central CDU to the RackCDUs™. The RackCDUs™ are passive.

As noted above, there are two pumps installed in parallel that circulate water between the central CDU and the fluid coolers located in the AHUs. One of these was operating during Asetek's site visit. The pumps have VFDs but, since they are controlled to differential pressure and the central CDU has a three-way valve, the flow rate is constant.

The fluid cooler fan speed is controlled to maintain a constant-supply water temperature. The modeling team found the setpoint temperature to be 68°F (20°C) during their site visit, and it was later raised to 75°F (24°C).

The fluid cooler performance is constrained by its approach to dry-bulb temperature (that is, the temperature difference between the cooling water and the dry-bulb temperature of the air entering the coil) and the evaporative cooler's approach to wet-bulb temperature (the temperature difference between the leaving dry-bulb temperature of the air from the evaporative cooler and the wet-bulb temperature of the ambient air entering the evaporative cooler). When in dry mode, the former applies; for example, the cooling water runs about 10°F (-12°C) above the ambient dry-bulb temperature. When in evaporative mode — if, for example, the ambient wet-bulb temperature is 65°F (18°C) and the approach temperature is 9°F (-13°C) — then the leaving dry-bulb temperature from the evaporative section would be 74°F and the cooling water would run at 84°F (29°C).

Observations

- The second floor had hot aisle containment rising above the racks. There was higher pressure on the contained hot aisle, due to server fans blowing into this aisle.
- The return air ducts were located on the cold side.
- The second floor needed a return plenum to significantly reduce mixing. Installation of a return air plenum was scheduled for 2019.
- The facilities staff indicated that the building had airflow issues that resulted in uneven cooling across the cluster and caused erratic performance and throttling across the cluster.

CHAPTER 2: Project Approach — Measurement and Verification Plan

Model-based Measurement and Verification Approach

Servers operated sporadically and generated less heat when idle. The day-to-day cooling requirements could consequently vary, depending on IT load. This, combined with the fact that the Asetek systems were installed in some new servers, means a traditional pre- and post-retrofit M&V analysis does not properly capture the Asetek system savings.

Instead, models were developed predicting the amount of power needed to cool each data center, based on the IT load and OAT. The savings were thus determined by comparing whether more heat was needed to cool the building when the Asetek liquid-cooled servers were operating or when the air-cooled servers were operating.

Models were developed individually for each building. Models A and C attempt to calculate energy use based on modeling equipment energy use and interactions. Models B and D are crude statistical models that correlate energy use with past IT load and temperatures. These models are represented as Options for each of the buildings.

Building 541

Building 451 housed the air-cooled Cab cluster that was then retrofitted with the Asetek liquid cooling solution.

Option A

During the modeling team's site visit, the members collected equipment data, discussed operation, and collected a variety of spot readings from equipment control screens. The team also received 10 days of pre- and post-retrofit trend data, including the following:

- Asetek RackCDU™ heat loads
- Cab IT kilowatts (kW)
- Remaining data center IT kW
- CHW pump kW
- Chiller kW
- kW and fan speed for all CRAH units
- Supply and return air temperature for all CRAH units

No data was available for the condenser water plant, so a constant performance of 0.58 kilowatts per ton (kW/ton) was assumed based on a recommendation provided by LLNL.

Model Development

The savings for the Asetek retrofit came 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 resulted in less heat being rejected to the CHW loop, reducing the energy consumption by the CHW and condenser water plants. The Asetek pumps slightly increased energy use. Only the condenser water plant efficiency was affected 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, steady-state operations were assumed.

The CRAH fan power baseline 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.

The chiller power trends could not be used to directly calculate CHW plant savings because of the trend data discrepancy described in the Data Collection section below. Instead, calculation of the CHW plant performance baseline (in kW/ton) was based on the trended chiller kW, pump kW, and cooling load (IT load plus CRAH fan power). Then the reduced fan power of the CRAH fans (the cooling load reduction) was multiplied by the combined CHW plant and condenser water plant performance to determine the cooling power savings.

The developed energy model was used to calculate the annual power usage effectiveness (PUE) for the data center and the annual partial PUE (pPUE) for various subsystems. The following was included:

- Data Center — Baseline and post-retrofit PUE
    - CRAHs
    - CHW plant
    - Condenser water plant
    - Total IT load
- Cab pPUE
    - CRAHs (only the energy associated with Cab)
    - CHW plant
    - Condenser water plant
    - Asetek pumps (CDU circulating pump)
    - Cab IT load
- Asetek w/ CHW pPUE
    - Asetek pumps (CDU circulating pump)
    - CHW plant (associated load)
    - Condenser water plant (associated load)
    - Cab IT load
- Asetek Only pPUE
    - Asetek pumps (CDU circulating pump)
    - Cab IT load

Option B

The building was equipped with the PI[1] database (a system provided by OSIsoft), LLNL's overarching system that acts as a backbone for the collection of data from different data sources and devices for recordkeeping and analysis. Based on the available data, the following data points (Table 6) were available to analyze the performance of the system.

Model Development

Based on engineering principles, mechanical power to cool the data center was a function of the cooling load and the outside weather conditions. The internal load was dominated by the IT load from various clusters, including Cab (a high-performance computing, or HPC, cluster cooled in part by Asetek technology). The IT power was also gathered for other clusters that were air-cooled. An initial analysis to understand the variables that affected the mechanical energy consumption Y-axis on Figure 3 indicated that the IT loads and OAT, X-axis, were significant.

Based on the analysis, the following variables were found to be significant to predict the mechanical power draw:

- Y is the total mechanical power in kW that includes both the CRAH units and the chiller loads identified as TB821 and TB1866, respectively, in Table 6
- X1 represents the IT load drawn from Cab and other air-cooled clusters, as shown by the TB1864 meter
- X2 represents the IT load drawn from air-cooled clusters not monitored by TB1864, as shown by the TB1865 meter
- OAT represents the outside air temperature in Fahrenheit

Table 6: Summary of 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	IT Load 2	TB1865 (kW, kWh)	Power draw by the IT equipment, partial
IT Load	Cab Load	Cab power (kW, kWh)	Power draw by the water-cooled Cab cluster (included in IT Load 1)
Mechanical Load	CRAH Load	TB821 (kW, kWh)	Power draw from all CRAH units; also feeds some pumps in the chiller room.

[1] PI is a real-time data historian application with a highly efficient time-series database that was structured and developed by OSIsoft. PI stands for process information.

Category	System	Variable	Description
Mechanical Load	Chiller Load	TB1866 (kW, kWh)	Chiller power; also supports the office area that is 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 in B451.
CDU Data (post-retrofit)	Heat captured by the water in each of the RackCDUs™	RackCDU™ Heat Load RackCDU™ Facility Temperature in RackCDU™ Facility Temperature Out RackCDU™ Facility Flow	Individual data from 16 different rack CDUs
Other	Weather	OAT (F)	
Other	Cluster/node data	Internal room, and server inlet temperatures	

Source: Asetek, Inc.

Figure 3: Power Relationships

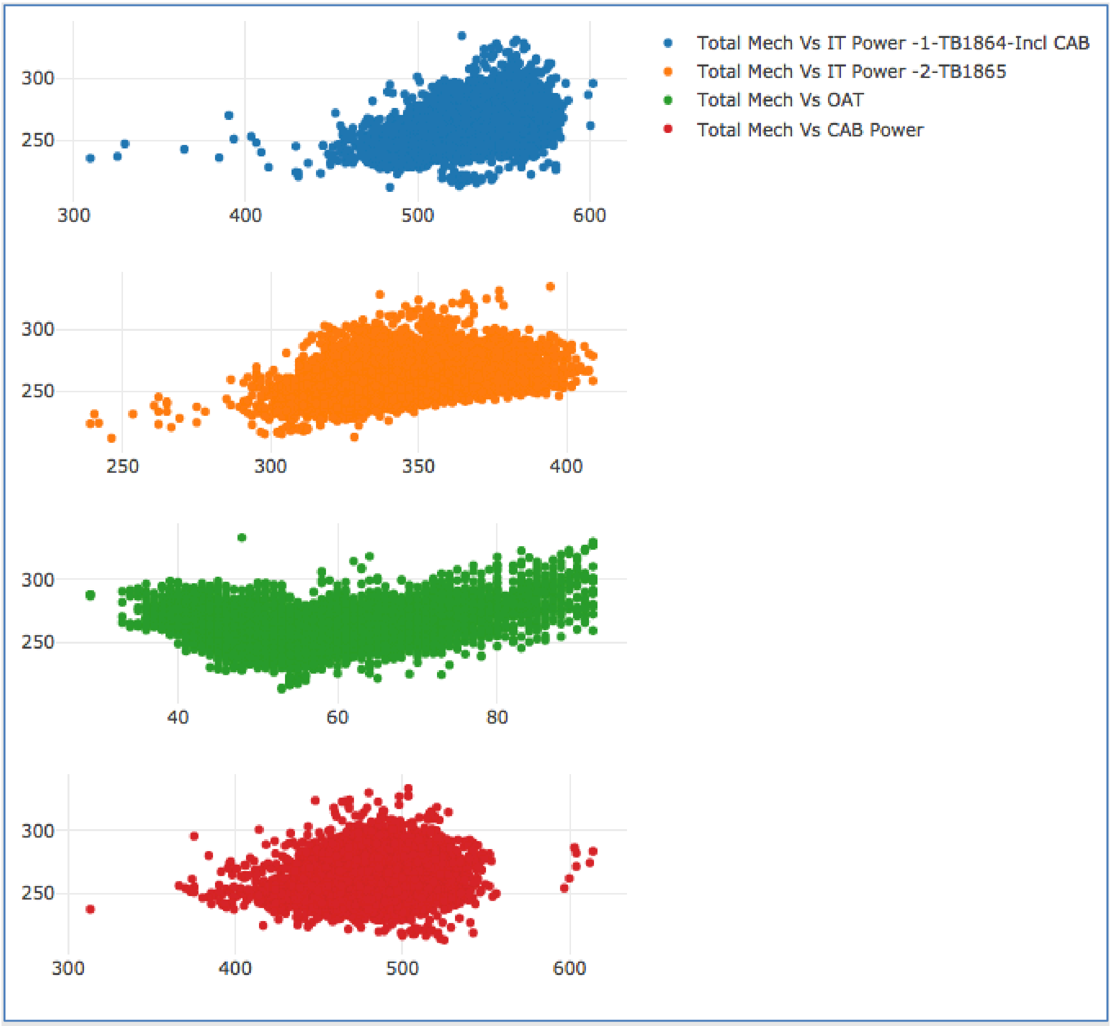

Plots showing the relation between total mechanical power (kW) on the Y-axis versus: (blue) IT Power 1-TB1864, including Cab (kW); (orange) IT Power 1-TB1865 (kW); (green) OAT (F); and (red) Cab Power (kW).

Source: Asetek, Inc.

Regression Model

Based on the regression analysis, the model (Equation 1, below) was developed to predict the mechanical power need for the computer racks for a given weather condition. 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 X1 (IT power 1), which includes the power draw by the Cab cluster, increases mechanical power to cool the cluster by 15 percent. On the other hand, a unit increase in X2 (IT power 2), which represents the air-cooled IT equipment, increases the mechanical power by around 30 percent. Table 7 summarizes the goodness of fit for the developed model.

Equation 1: Equation for Regression Analysis

$$Y = 43.1218 + 0.1525 * X1 + 0.2915 * X2 + 0.5991 * OAT$$

Table 7: Summary of Model Fit Characteristics

Statistic	Description	Model
Root mean squared error (RMSE)	An indicator of the scatter, or random variability, in the data, and hence an average of how much an actual Y value differs from the predicted Y value	11.45
Coefficient of variation RMSE (CVRMSE)	Nondimensional metric that normalizes RMSE by the average Y value that describes how well the model fits the data	4.4 percent
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 percent
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 percent

Source: Asetek, Inc.

One of the key metrics is the coefficient of determination, R_2, which measures the extent to which variations in the dependent variable Y can be explained by the regression model. The possible range for R_2 is between 0 and 1. A value of 0 indicates that none of the variation can be explained by the model and, therefore, the model provides no guidance in understanding the variations in Y using the selected independent variables; a value of 1 means that the model explains 100 percent of the variations in Y. For a good model, the general criterion for R_2 is to be more than 70 percent (EVO, 2012). For the model, this metric was calculated at 74.3 percent, which indicates that 74 percent of the variation in mechanical power can be explained by the predictor variables X1, X2, and OAT.

The 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.

The coefficient of variation RMSE (CVRMSE) is the RMSE normalized by the average Y value. Normalizing the RMSE makes this a unitless number that describes how well the model fits the data. The general criterion for CVRMSE is that it should be less than 25 percent. It is not affected by the degree of dependence between the independent and the dependent variables, making it more informative than R-squared for situations where the dependence is relatively low.

The NBE is simply the 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. If the net determination bias equals zero, then there is no bias. The general criterion for the NBE is that it should be less than 0.5 percent. The developed baseline model was used to predict the load consumed by the mechanical system that included both the chiller and the CRAH systems for the base year (shown in dark brown in Figure 4) and compared with the actual mechanical load.

Figure 4: Plot Showing Model Predicted and Actual Mechanical Baseline Power

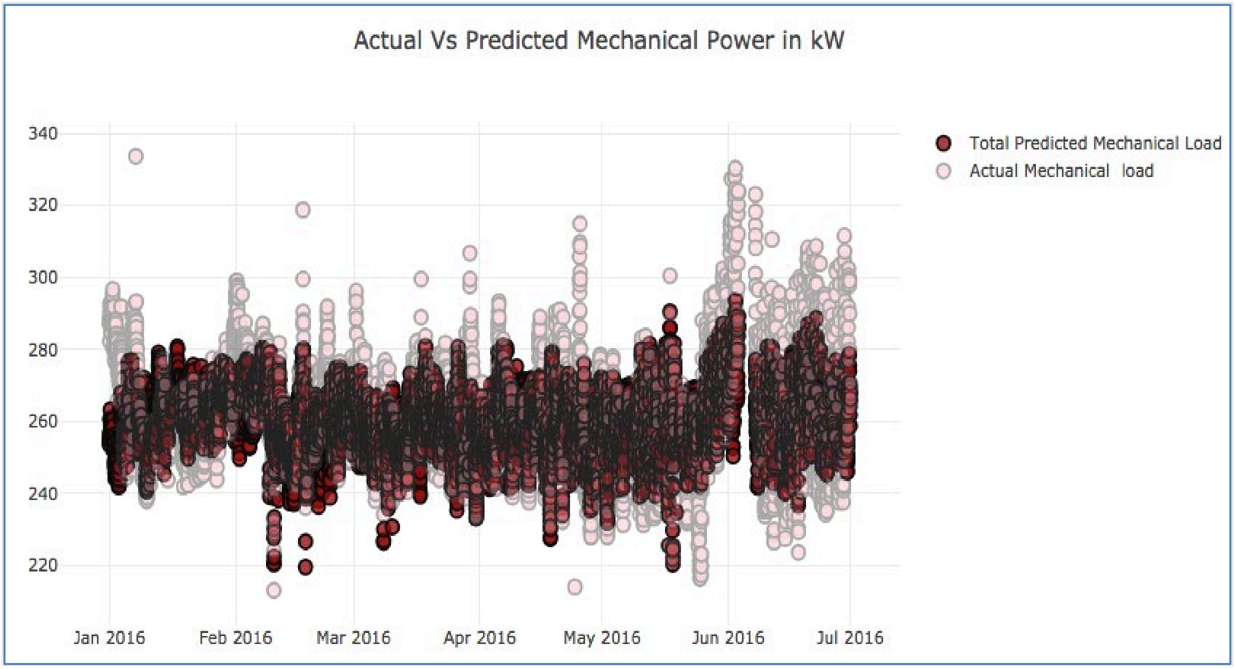

Source: Asetek, Inc.

Savings Analysis

The developed baseline model was used to project the load consumed by the mechanical system that included 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 include the IT power and OAT were used to project baseline load for the overall mechanical system for each of the 30-minute intervals. This projected baseline mechanical load (blue line in Figure 5) was assumed to be the load as if no retrofit was implemented. This predicted mechanical baseline load was compared with the corresponding actual mechanical load from 2017 (orange line in Figure 5) to assess the effect of the Asetek retrofit.

Figure 5: Plot Showing Model Predicted and Actual Mechanical Baseline Power

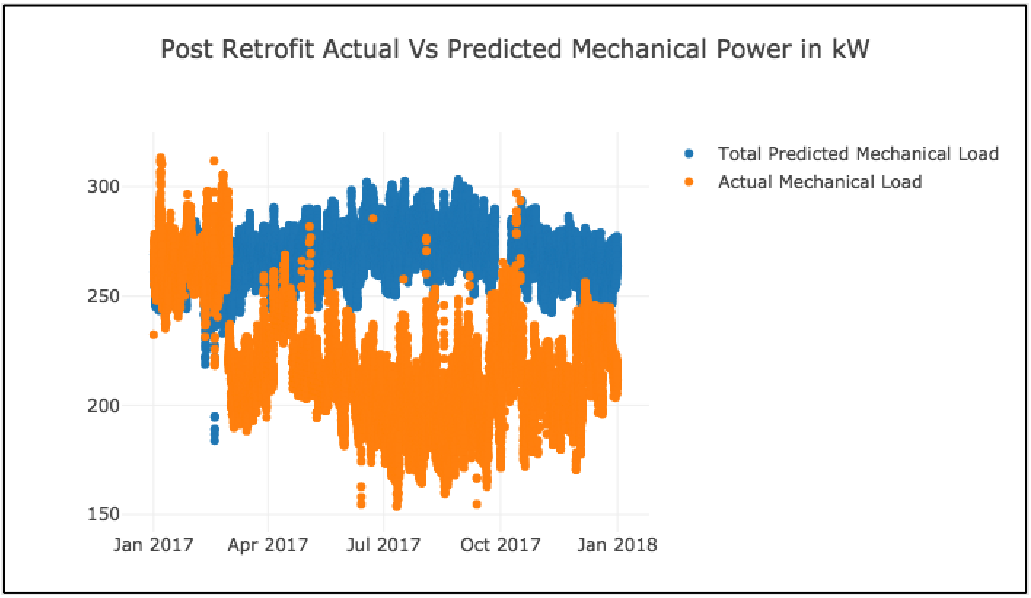

Source: Asetek, Inc.

Building 654

Building 654 simultaneously housed the air-cooled cluster (QUARTZ) and the liquid-cooled cluster (TOPAZ)

Option C

This approach involved developing physics-based engineering models to estimate the baseline and post-retrofit energy consumption for the measure. These models were based on weather-based bin simulations for both airside and waterside systems using mechanical schedules, site visit notes, and sequences of operations provided by LLNL. These models were calibrated based on spot measurements taken during the site visit and power trends provided for IT power, total mechanical power, and CDU load.

Data Collection

During the team's site visit, equipment data was collected, operation parameters were discussed, and a variety of spot readings from equipment control screens were collected. Readings from the VFDs for the supply fans, fluid cooler pumps, and fluid cooler fan were collected to determine speed and power. The exhaust fan VFDs were not accessible. The team also discussed operation with the facilities staff later to acquire additional spot readings of equipment power during warmer conditions.

The team received trend data that included the following:

- Power draw for Quartz
- Power draw for Topaz
- Infrastructure power draw (everything else)

No other trend data was available for B654.

Model Development

The Asetek team developed calculations using a 2°F (-17°C) bin simulation with TMY3 data for Livermore, California. The simulation was built using spot measurements and the team's understanding of the mechanical system sequences of operations. The flow rate of the supply fans, CDU pump, and fluid cooler pumps were constant, so the power draw was based on spot measurements. The Asetek rack pumps were also constant speed and were all assumed to have a draw of 4 W, as stated by Asetek. Since they were powered by the IT power supplies, these pumps were counted as IT load.

An assumption was made that the exhaust fan speed would vary proportionally to outside air flow rate. The team calculated the outside air flow rate based on OAT, assuming the supply air temperature (SAT) and return air temperature (RAT) remained constant, as seen on site. The power at reduced speeds was determined by using the power of the exhaust fan at full speed from a spot measurement and an assumed VFD coefficient of 2.4 (that is, the input power scaling with the 2.4 power of the speed).

The fluid cooler fan power was based on a spot measurement at an outside air wet-bulb temperature of 47°F (8°C) and one at an outside air wet-bulb temperature of 68°F (20°C) (when the fans were operating at 100 percent speed). The approach temperature of 10°F (-12°C) was measured when the fans were operating at 100 percent speed. The fan speed was determined by assuming a constant approach to wet-bulb and reducing fan speed proportionally when the load was met. It was found to be consistent with the reduced fluid cooler fan power Asetek measured at 47°F (8°C).

Lighting power was determined based on a typical lighting power density of 1.2 watts per square foot (W/sf) and assumed distribution losses of 3 percent of the IT power. Distribution losses were increased in this model from a default of 2 percent to better match trend data, and the team believes it to be reasonable.

The final steps were to:

1. Develop a model that correlated the mechanical load as a function of OAT, both the IT loads (air-cooled Quartz and water-cooled Topaz), and other possible parameters.
2. Select the model with significant independent variables that could satisfactorily explain the variation in the mechanical power.
3. Validate the model to ensure that goodness of fit conditions were being met while ensuring that the assumptions of the regression model were true.
4. Finalize the model that predicted the kW drawn by the mechanical equipment, including AHUs, fluid coolers, and other equipment that were needed to provide both air- and water-cooling to remove the heat generated by the IT equipment.
5. The final steps led to the following equation:

Equation 2:

$$\mathrm{kWmech} = \mathrm{a} + \mathrm{b1(X1)} + \mathrm{b2(X2)} + \mathrm{b3(X3)} + \mathrm{b4(X4)} + \mathrm{b5(X5)}$$

where X1, X2, X3, X4, and X5 are independent variables that can adequately explain the variations in the mechanical power (kWmech) drawn by the cooling equipment (for example, Topaz IT power, Quartz IT power, OAT) while b1, b2, b3, b4 and b5 are their corresponding coefficients calculated through analysis.

Option D

Data Collection

The building was equipped with the PI database (a system provided by OSIsoft), LLNL's over-arching 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 8) were available to analyze the performance of the system.

Table 8: Summary of Data Variables Available for Building 654

Category	System	Variable	Description
IT Load	Quartz IT load	TB1938 (kW, kWh)	Power draw by the Quartz cluster
IT Load	Topaz IT load	TB1937 (kW, kWh)	Power draw by the Topaz cluster
Mechanical Load	Mechanical load		Power draw from AHUs, including all the supply fans, exhaust fans, and fluid cooler fans
CDU Data	Heat captured by the water in each of the Rack CDUs	RackCDU™ Heat Load RackCDU™ Facility Temperature in RackCDU™ Facility Temperature Out RackCDU™ Facility Flow	Individual data from 12 different rack CDUs
Other	Weather	OAT (°F)	
Other	Cluster/node data	Internal room, and server inlet temperatures	

Source: Asetek, Inc.

Both clusters, Topaz and Quartz, were commissioned in December 2017. However, the data systems were installed and equipped to log and trend data from early January of 2018. The available data were from February 3, 2018, to May 29, 2018, and were sampled at five-minute intervals (Figure 6).

Figure 6: IT and Mechanical Loads With Outside Temperature at Building 654

Source: Asetek, Inc.

However, the CDU was available for only a month starting in early March 2018. Benchmarking runs were performed on Topaz on March 15 from noon to 5:30 p.m. and on Quartz on February 18 from 10 am to 4pm.

The data available from the PI system included variables listed, along with the time stamp. The data used for analysis were sampled at 30-minute intervals. These plots showed some of the anomalies in the data that included outliers and possible erroneous values.

Model Development

Based on engineering principles, mechanical power needed to cool the data center scales with the cooling load and the outside weather conditions. The internal load was dominated by the IT load from clusters, including those that were water-cooled clusters using Asetek technology. An initial analysis to understand the variables that affected the mechanical energy consumption (Y axis on Figure 7) indicated that the IT loads and OAT (X axis) were significant.

Figure 7: Power Relationships

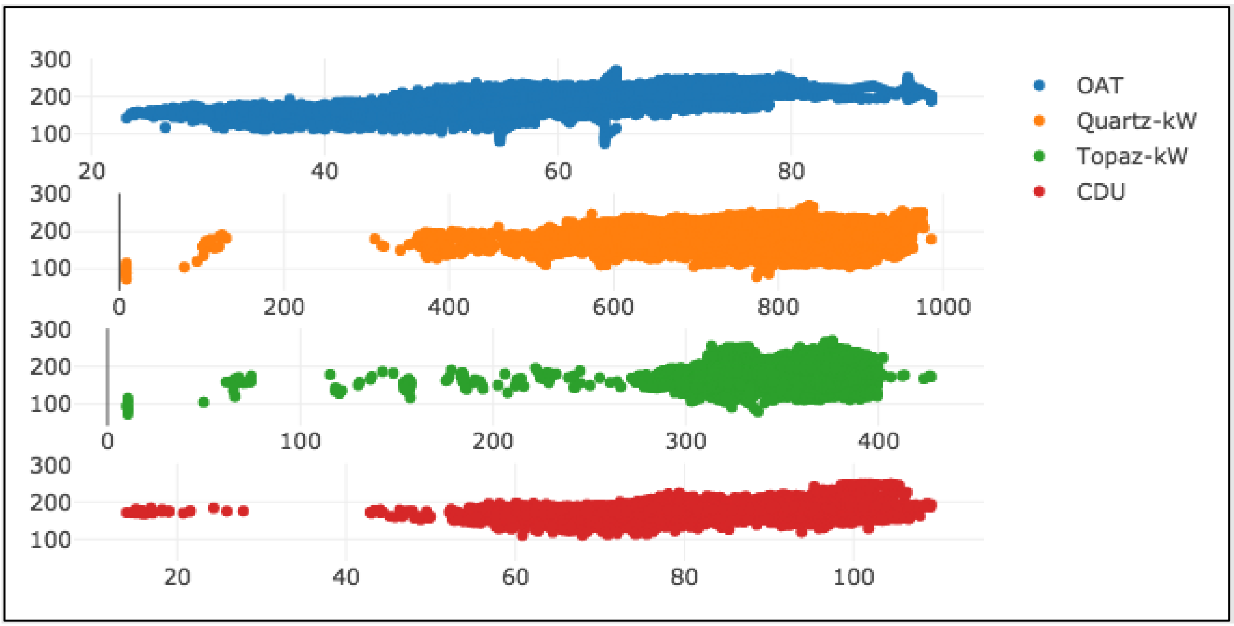

Plots showing the relation between total mechanical power (kW) on the Y axis versus: (blue) OAT; (orange) Quartz power; (green) Topaz power; and (red) heat removed by CDU.

Source: Asetek, Inc.

Based on the analysis, the following variables were found to be significant to predict the mechanical power draw:

- Y is the total mechanical power in kW that includes both the load drawn by AHUs, supply fans, exhaust fans, and fluid cooler fans and power drawn by Asetek and CDU pumps
- X1 represents the Topaz IT load drawn (kW)
- X2 represents the Quartz IT load drawn (kW)
- OAT represents the outside air temperature in °F

Regression Model

Based on the regression analysis, the following model was developed to predict the mechanical power need for the computer racks for a given weather condition. The regression coefficients shown in Equation 3 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 X1 (Topaz) increases mechanical power to cool the cluster by 3 percent. On the other hand, a unit increase in X2 (Quartz), the completely air-cooled cluster, increases the mechanical power by around 0.5 percent.

Equation 3:

$$\mathbf{Y = 79.59 + 0.0305 * X1 + 0.0053 * X2 + 1.494 * OAT}$$

Table 9 indicates the summary to assess goodness of fit for the developed model. The coefficient of determination, R_2, was calculated to be 57 percent, which indicates that 57 percent of the variation in mechanical power could be explained by the predictor variables — X1, X2, and OAT — and does not meet the requirements for a good model.

Table 9: Summary of Model Fit Characteristics

Statistic	Description	Model
Root mean squared error (RMSE)	An indicator of the scatter, or random variability, in the data, and hence an average of how much an actual Y value differs from the predicted Y value	9.6
CVRMSE	Nondimensional metric that normalizes RMSE by the average Y value that describes how well the model fits the data	5.5 percent
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	0.1 percent
Coefficient of determination (R_2)	Measures the extent to which variations in the dependent variable Y can be explained by the regression model	56.6 percent

Source: Asetek, Inc.

The R_2 test should be used only as an initial check. Models should not be rejected or accepted solely on the basis of R_2. Regression models with low R_2 values can be perfectly good models for several reasons. Some data sets have an inherently greater amount of unexplainable variation. In these areas, R_2 values are bound to be lower. In cases where the R_2 value is low but the independent variables are statistically significant (where the T-statistic is greater than 2.0 or the p-value is less than 0.05), one can still draw important conclusions about the relationships between the variables. Statistically significant coefficients continue to represent the mean change in the dependent variable given a one-unit shift in the independent variable.

The 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. The RMSE of 9.6 translates to a CVRMSE of 5.5 percent and meets the general criterion for CVRMSE to be less than 25 percent. It is not affected by the degree of dependence between the independent and the dependent variables, making it more informative than R-squared for situations where the dependence is relatively low. The NBE was calculated to be 0.1 percent, which is less than the 0.5 percent requirement for a good model.

The developed model was used to predict the load consumed by the mechanical system for the study period in 2018, given the IT loads from Topaz and Quartz, along with the OAT data. This predicted load (shown in orange in Figure 8) was compared with the actual mechanical load to assess the overall fit of the model.

Figure 8: Plot Showing Model Predicted and Actual Mechanical Baseline Power for Building 654

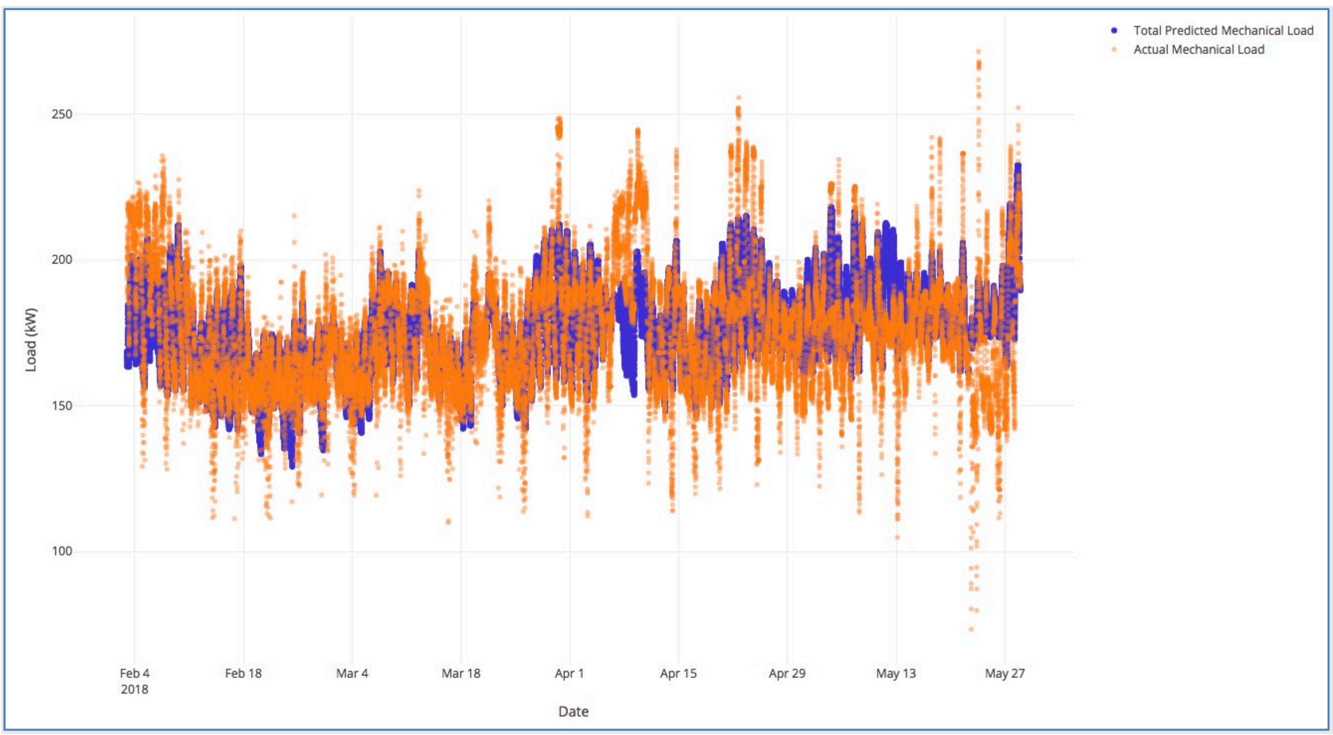

Source: Asetek, Inc.

CHAPTER 3: Project Results

Option A Results

The energy model developed for Building 451 was used to create energy balances of the whole data center and two different subsystems. The energy balance in Figure 9 shows the calculated baseline annual energy use percentages for all systems in the data center. Figure 10 shows the same systems after the Asetek system was installed.

Figure 9: Baseline Full Data Center Energy Balance (kWh)

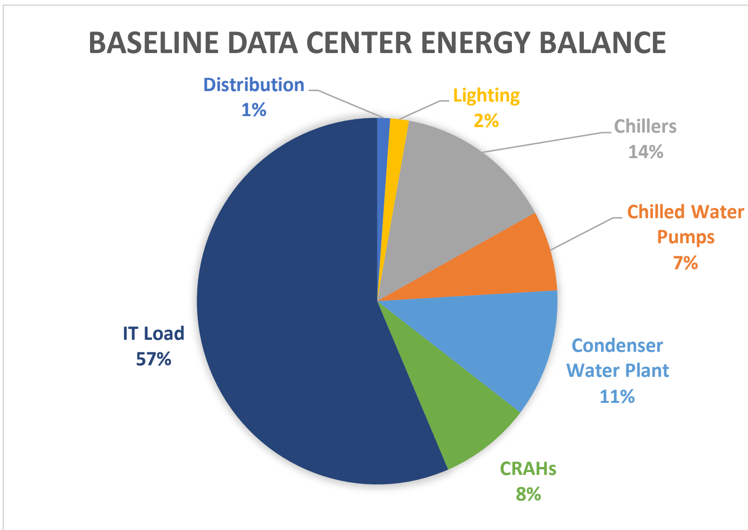

Source: Asetek, Inc.

Figure 10: Post-retrofit Full Data Center Energy Balance (kWh)

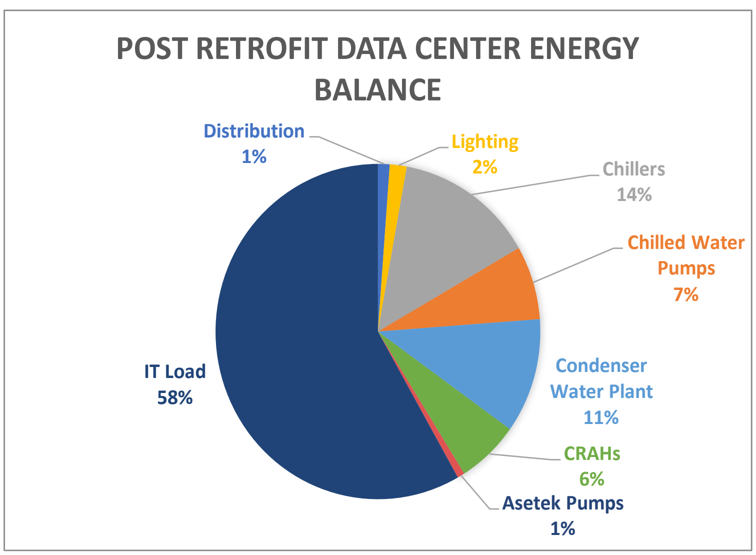

Source: Asetek, Inc.

Figure 11 shows the energy balance of the Cab IT load and the associated systems serving it. According to trends, the Asetek system removed 33 percent of the heat generated by Cab, with the rest of the load remaining on the CRAH units.

Figure 11: Post-retrofit Cab (kWh)

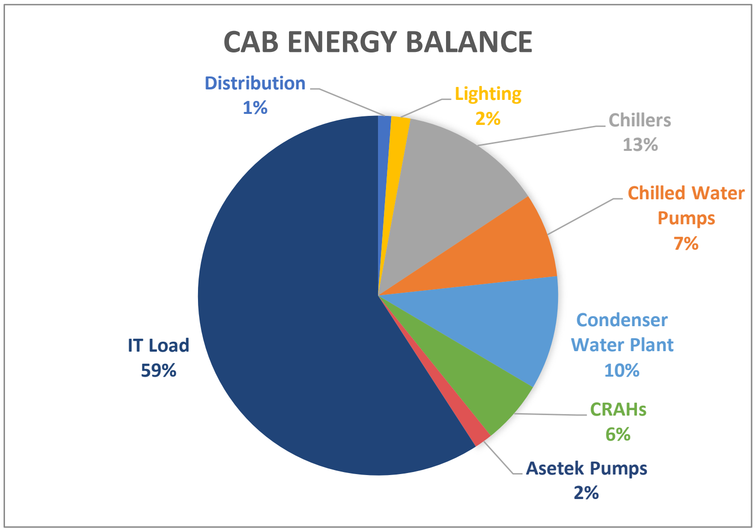

Source: Asetek, Inc.

Table 10 summarizes the PUE and pPUEs defined for the proposed upgrade.

Table 10: PUE and pPUE Summary for B451

Baseline: Data Center PUE	Data Center PUE	Post-retrofit: Cab pPUE	Post-retrofit: Asetek With CHW pPUE	Post-retrofit: Asetek Only pPUE
1.64	1.59	1.52	1.41	1.01

Source: Asetek, Inc.

The total overall data center PUE reduced from 1.64 to 1.59 as a result of the retrofit. The Asetek pumps within the servers were part of the IT load. In a properly configured server, the internal fans would slow down, as they do not need to move air through the CPU air heat sinks, and there would be a net energy savings at the server level. In the case of Cab, the server fans were set to a constant speed and LLNL was unable to slow them; thus the server level energy savings was not demonstrated. Here the pPUE shows that the Asetek system was very efficient at removing heat from the servers it served, as this pPUE is calculated to be 1.01. However, since the Asetek system rejected heat to the building CHW system, once the CHW and condenser water plant efficiencies were considered, the pPUE was much higher and this "Asetek w/CHW" pPUE was calculated to be 1.41. Based on the data, LBNL determined that the combined CHW and condenser water plant efficiency was relatively low. Rejecting heat directly to an efficient condenser water loop would reduce the PUE significantly.

The Cab pPUE (1.52) isolates Cab and its associated cooling systems, including both CRAH fans and the Asetek system. It is Asetek's best estimate for what the PUE of the data center would be if all servers were cooled by Asetek. Asetek expected this reduction to have been much more significant if the Asetek system had not been required to reject heat to the CHW system.

Option B Results

The actual mechanical load for the post-retrofit conditions was unusually high for the first two months of 2017 and no plausible explanation was obtained from the implementation team or the site team members. These data were retained and not excluded from the analysis.

The 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]
>
> PUEPR = [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 = the load from office rooms, lighting, and plug loads, assumed not to 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 [British thermal units (Btus)/kWh]/12000 [Btu/ton-hr)] * Condenser Water kW/ton1
>
> 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[2]

[2] Estimated by LLNL as 0.58 and assumed to remain constant and not affected by Asetek.

Figures 12 compares the baseline PUE (PUE_{BL}) with post-retrofit PUE (PUE_{PR}) conditions.

Figure 12: Predicted Baseline and Actual PUE for Post-retrofit Period

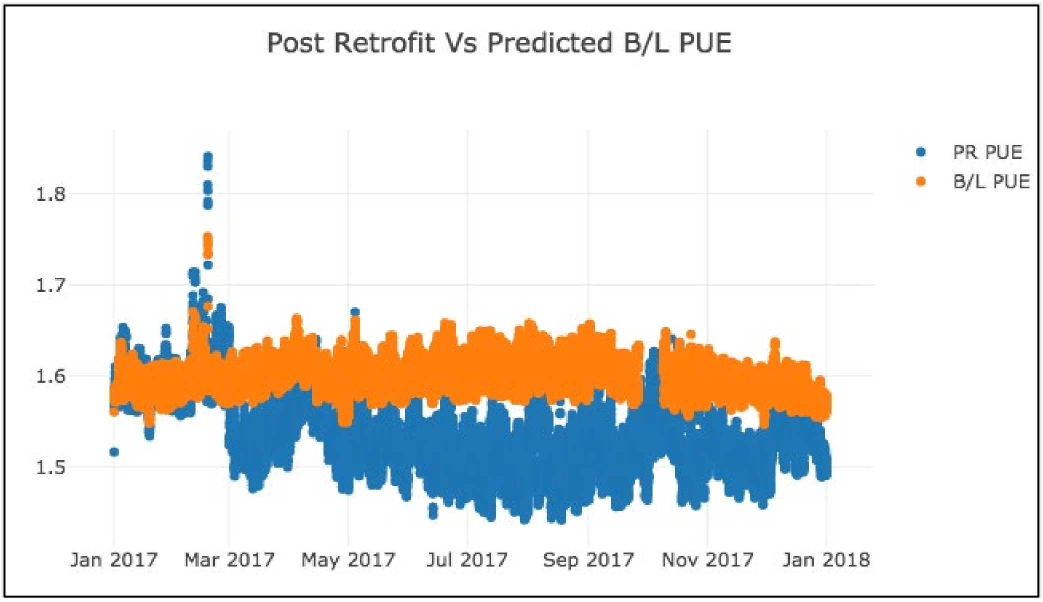

Source: Asetek, Inc.

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 was included in this analysis; when included, it would increase the post-retrofit PUE slightly.

The heat gathered by all the Asetek RackCDUs™ was compared with the Cab power for a 10-day period in February 2018. This heat gathered by CDUs was found to be 37 percent 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 and Greenberg, 2014) found that heat capture rates by the Asetek system ranged from about 47 percent to 63 percent (depending on IT load) at the same cooling water supply temperature to the rack CDUs as in the B451 case. However, direct comparisons are complicated by the differences between the LBNL study and B451, as shown in Table 11.

Table 11: Comparison of Conditions at Building 451 and the LBNL 2014 Study

Parameter	LBNL 2014	LLNL B451
Water-cooled components	Processors and memory	Processors only
Supply water temperature	59, 68, 77, 86, 104, 113°F (15, 20, 25, 30, 40, and 45° C)	79°F (26° C) average
Water flow rate per rack	4.9 gpm	3.9 gpm average
Server power	120, 270, and 430 W each	270 W average
Room air temperature	82°F (28° C)	<72°F (22° C) average

Source: Asetek, Inc.

The 2014 configuration included water-cooled memory and used higher water flow rates than the LLNL study, both of which would result in higher capture rates at LBNL. The room temperature was higher in 2014 than for the LLNL study (the 72°F [22°C] is the average of the CRAH return temperatures, with the CRAH supply temperatures averaging about 61°F [16°C]), also resulting in higher capture to water in 2014. The HPC IT equipment and software used for loading the IT equipment was also 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 et al., 2014), the National Renewable Energy Laboratory (NREL) found 48 percent of the IT load rejected to water with a system using water-cooled CPUs only (not memory chips), 1.1 gpm of 63°F (17°C) cooling water supplied to the RackCDU™, and 68°F (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 percent in the LBNL test under very similar water temperatures, load, and flow, but with the memory chips also water-cooled.

Option C Results

The energy balance from the model (Figure 13) shows the calculated annual energy use percentages for all systems in Building 654. The "IT Load" in the figure is the combination of Quartz and Topaz energy consumption.

Figure 13: Whole Building Energy Balance (kWh)

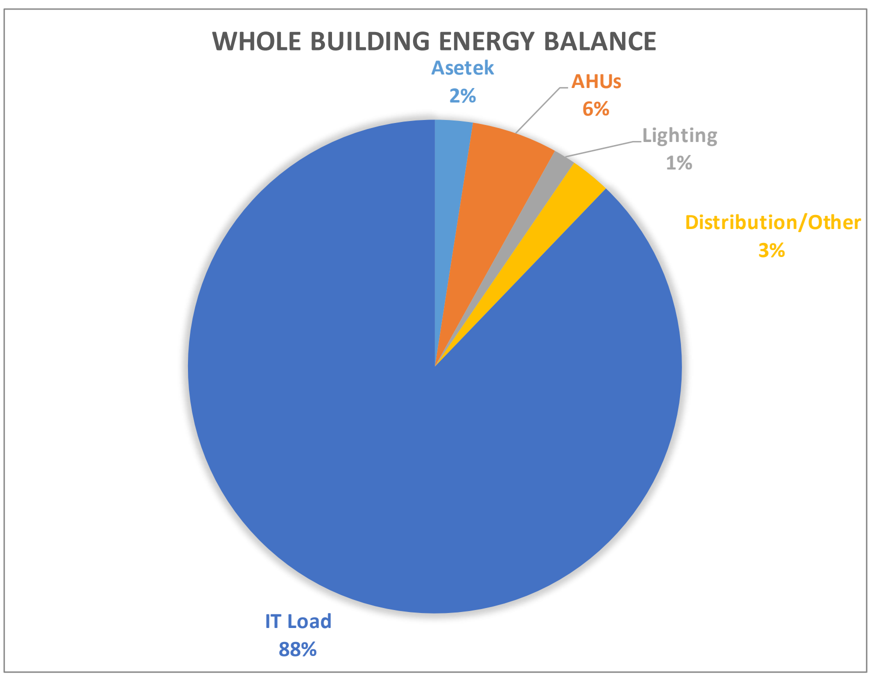

Source: Asetek, Inc.

The AHU energy balance shown in Figure 14 shows the annual energy use percentages for the two AHUs and the IT load they serve. The IT load was calculated by subtracting the heat removed by the Asetek system from the total IT load.

Figure 14: Energy Balance for AHUs and IT Load Served (kWh)

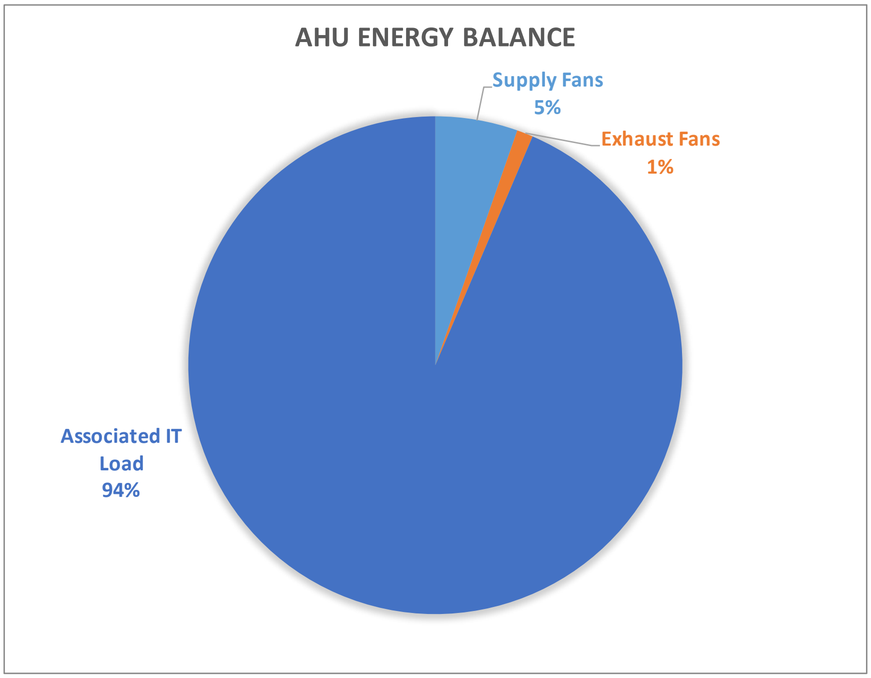

Source: Asetek, Inc.

This third energy balance (Figure 15) shows the annual energy use percentages of the Asetek system and the associated IT cooling load. Asetek determined the average cooling load from RackCDU™ trends. Based on these CDU trends and power trends for Topaz, it appeared that the Asetek system removed 23 percent of the heat generated by Topaz.

Figure 15: Energy Balance for the Asetek System and IT Load Served (kWh)

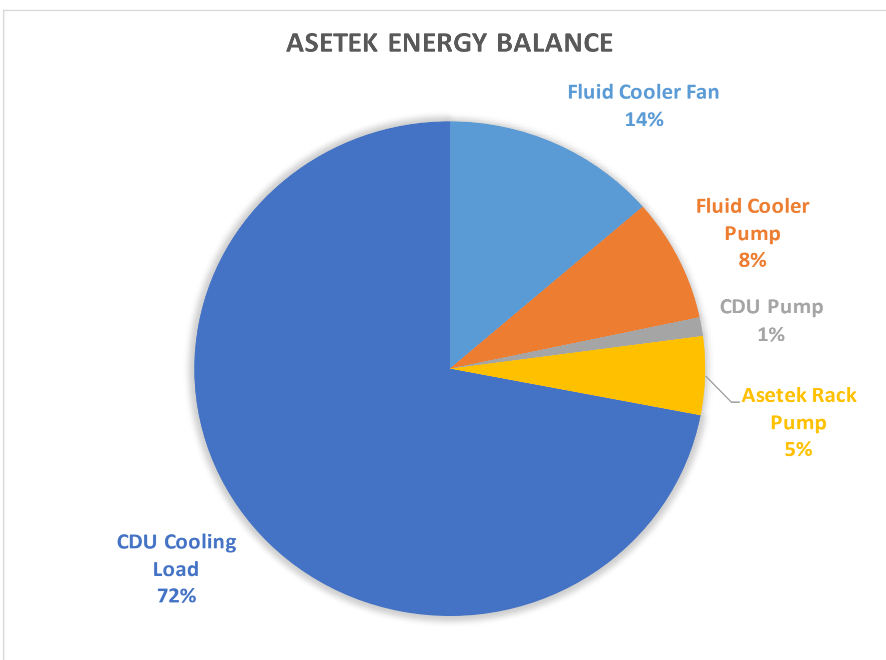

Source: Asetek, Inc.

The energy model the team developed was used to calculate the annual PUE for the building and the annual partial PUE (pPUE) for various subsystems (Table 12). Here the pPUE is the PUE of one individual system, using only the IT load it serves. The following was included:

- AHU pPUE
    - Supply fans
    - Exhaust fans
    - Total IT load not served by Asetek (partial Topaz and all of Quartz)
- Asetek PUE
    - Fluid cooler fan
    - Pumps for fluid cooler, CDU, and racks
    - IT load served by Asetek system (cooling provided by the RackCDUs)
- Total PUE
    - AHUs and Asetek system
    - Lighting and distribution
    - Topaz and Quartz

Table 12: Calculated Partial PUEs and Total Building 654 PUE

AHU pPUE	Asetek pPUE	Total PUE
1.07	1.32	1.13

Source: Asetek, Inc.

The resulting pPUE for the Asetek system was found to be surprisingly high. Using similar systems as a reference, the pPUE was expected to easily fall below 1.2, and, in general, water cooling is expected to be more efficient (lower pPUE) than air cooling (Hughes, 2015). This difference was well above the expected margin of error for the team's calculations. More investigation would be required to definitively resolve this difference. The reported heat load for the Asetek system was only 23 percent of the Topaz power consumption, so it was possible that the reported heat load for the Asetek system was significantly lower than reality. However, the total heat load rejected, according to the RackCDU™ meters, agreed within 10 percent of the total heat rejected by the central CDU (107 kW versus 118 kW), which suggested that the heat rejection numbers, and capture fraction, were not far off.

The high pPUE for the Asetek portion of Topaz was due to a combination of low heat capture fraction and the inability of the water-cooling system to turn down efficiently at low loads. Regarding the heat capture fraction, the 81°F (27°C) supply water temperature at B654 would have resulted in about a 44 percent to 62 percent capture, depending on server load, in the LBNL configuration (Coles and Greenberg, 2014). In the LBNL study, Asetek heat sinks were included on both the processor and the memory chips. With a comparable supply and return temperature, the LBNL configuration would have performed in the 30 percent to 60 percent range, again depending on server load. Regarding the turn-down capability of the facility heat-rejection system, the flow to the central CDU was 87 gpm, for a pump designed for 800 gpm, but, due to pump speed control on differential pressure and the use of a three-way control valve at the CDU, much higher pumping power was used than would be under a different control strategy. Likewise, at the fluid cooler, modifying the control scheme was likely to make the system significantly more efficient.

Also note that the PUE was calculated to be 1.17 for the full building using the provided trend data alone. This was slightly different from the 1.13 shown above; however, further calibration would not be appropriate without more detailed trends.

Option D Results

This predicted baseline mechanical load for both hypothetical cases of the entire IT load provided by either Topaz or Quartz was calculated using the model. These numbers were, in turn, used to calculate the PUE for both cases for a given identical set of weather conditions. The PUE was calculated as:

PUE_Q = [Total IT Load by Quartz + Total Predicted Mechanical Load for Quartz]/[Total IT Load by Quartz]

PUE_T = [Total IT Load by Topaz+ Total Predicted Mechanical Load for Topaz]/[Total IT Load by Topaz]

Total IT Load by Quartz and Total IT Load by Topaz were assumed to be the same, which is the sum total of the load from Quartz and Topaz gathered from the meter data.

The PUE for the current configuration, where Quartz and Topaz were 70 percent and 30 percent, respectively, of the total IT load, was calculated to be 1.15. The PUE (PUE_Q) for a hypothetical case, where Quartz was 100 percent of the load, was found to be 1.14, which was lower than the PUE (PUE_T) of 1.17 for the other hypothetical case, where Topaz was 100 percent of the load, assuming the weather conditions were the same.

The heat gathered by all the Asetek RackCDUs™ was compared with the Topaz power for a one-month period (March 5 through April 5, 2018). This heat gathered by CDUs was found to be 23 percent of the overall Topaz IT power, which was considerably lower than what was found in previous research. For example, a study done in 2014 by LBNL (Coles and Greenberg, 2014) found that heat capture rates by the Asetek system ranged from about 44 percent to 62 percent (depending on IT load) at the same cooling water supply temperature to the RackCDUs™ as in the B654 case. However, direct comparisons are complicated by the differences between the LBNL study and B451, as shown in Table 13. In the Topaz machine at B654, there were significant non-CPU loads that were not liquid-cooled, contributing to a lower heat capture rate.

Table 13: Comparison of Conditions at Building 654 and the LBNL 2014 Study

Parameter	LBNL 2014	LLNL B654
Water-cooled components	Processors and memory	Processors only
Supply water temperature	59, 68, 77, 86, 104, 113°F (15, 20, 25, 30, 40, and 45° C)	81°F (27° C) average
Water flow rate per rack	4.9 gpm	7.0 gpm average

Source: Asetek, Inc.

The 2014 configuration included water-cooled memory and used higher water flow rates than at LLNL, both of which would result in higher capture rates at LBNL. 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 was known to speculate on which direction these differences would affect the capture rate. In another study (Sickinger et al., 2014), the NREL found 48 percent of the IT load rejected to water with a system using water-cooled CPUs only (not memory chips), 1.1 gpm of 63°F (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 percent in the LBNL test under very similar water temperatures, load, and flow, but with the memory chips also water-cooled.

Computational Performance Analysis

Task Objective

The goal of this analysis was to quantify the computational performance improvements of the second site, that is, building 654 of LLNL. However, given the challenges with timeline and location changes, Asetek was using data from analysis on Building 451.

Data Collection

Data was collected as described in Task 5.1, as well as in the report published by Marathe et al. (2017) at LLNL.

Results

Rack Temperature Variation

Asetek evaluated CPU and rack temperatures for the Cabcluster in Building 451. Asetek observed CPU and rack temperature variations before and after retrofitting the cluster with Asetek's liquid cooling solution. Asetek found that, pre-retrofit, the ambient rack temperature on the hot and cold side of the rack could get as high as 950°F (510°C), with significant swings in rack temperature between 70°F (21°C) and 95°F (35°C) across racks. After the system was retrofitted with liquid cooling solution, Asetek observed a significant improvement (decrease) in rack temperature variation, and a much lower rack temperature of 40°F (4°C) overall, as indicated in the chart shown in Figure 16.

Process Temperature Performance

The rack-level erratic temperature swings observed in the air-cooled racks were also observed at the processor level. The processor temperatures got as high as 167°F (75°C) and fluctuated between 122°F (50°C) and 167°F (75°C) across different processors. Liquid cooling resulted in lower temperatures and very minimal temperature fluctuation across nodes. The processor temperature remained between 113°F (45°C) and 122°F (50°C), as indicated in Figure 17. Note: P1 is the upstream processor (CPU) and P0 is the downstream processor. The down-stream processor received pre-heated air from the upstream processor and thus operated at a higher temperature.

Figure 16: Comparing Rack Temperature Between Liquid-Cooling System and Air-Cooled System at LLNL

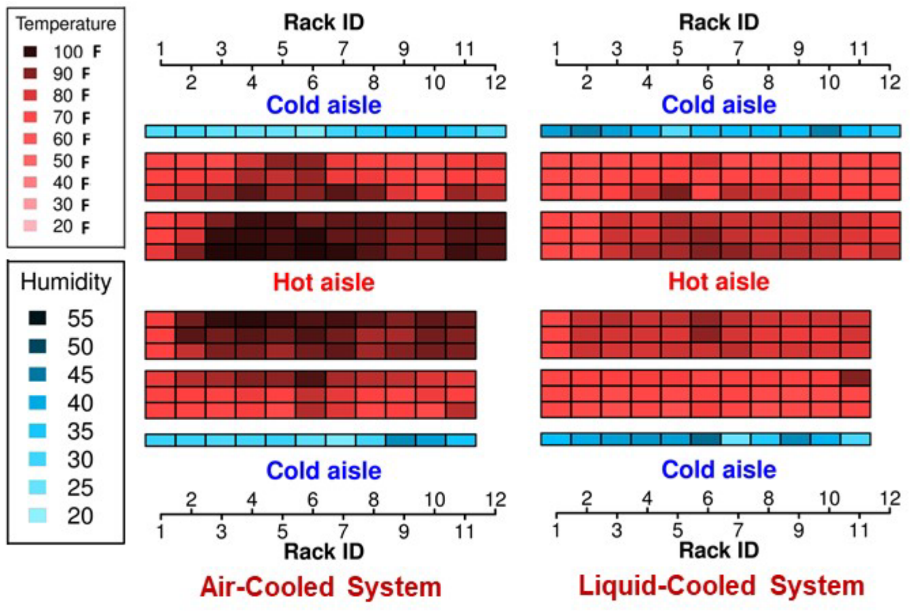

Source: Asetek, Inc.

Figure 17: IPC and Temperature of Liquid Cooling and Air Cooling

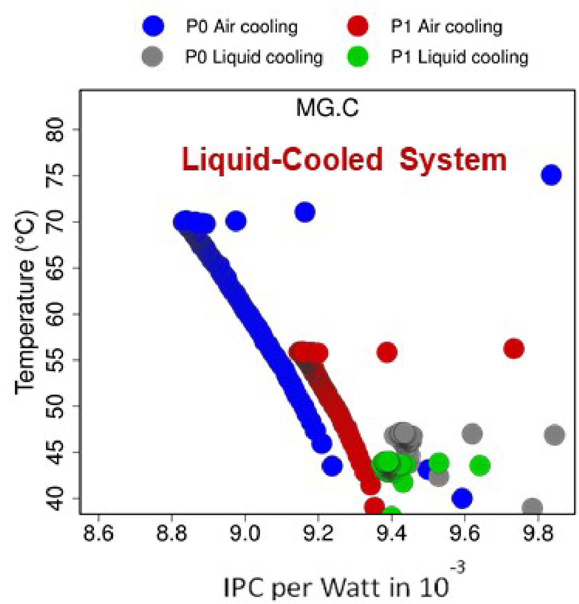

Source: Asetek, Inc.

Lower temperature operation and minimal fluctuation resulted in improved processor power efficiency. Liquid cooling minimized the difference in processor temperatures, overall operating temperatures, and a reduction in processor leakage power. The lower temperature enabled by liquid cooling also improved processor power efficiency, increasing efficiency — instructions per cycle per watt (IPC/W) — by 7 percent.

As shown in Figure 18, at 113°F (45°C), the CPU power draw was 102 W; however, as processor temperature increased, the average power draw increased to 108 W. This amounted to a power leakage of 6 W per processor. The work done by Osman Sarood et al. (2013) at the University of Illinois Urbana Champaign found that, for every 10-degree increase in processor temperature, there is a doubling of the soft error rate. The soft error rate was not measured for this experiment. However, using the work cited above, Asetek could extrapolate that, as CPU temperature increases, the soft error rate also increases. Hence, with the improved performance promised by liquid cooling, the soft error rate of the liquid-cooled system would be one-quarter of the air-cooled system (Table 14).

Figure 18: Power and Temperature Performance Comparison of Liquid Cooling Versus Air Cooling

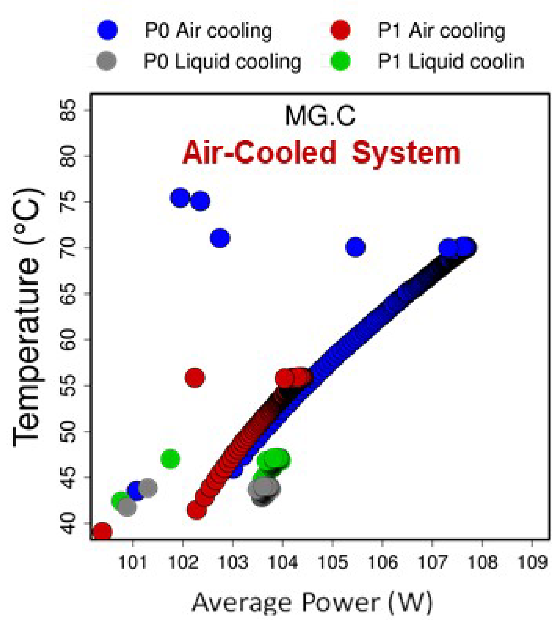

Source: Asetek, Inc.

Table 14: Tabulated Matrix of Performance Parameters

Parameter	Air-Cooled	Liquid-Cooled
Soft Error Rate	High	Low (A quarter of air-cooled)
Rack Temperature	High (95°F [35°C])	Low (75°F [24°C])
Power Efficiency	Low	High
Rack Temperature Fluctuation	High (70°F–95°F [21°C–35°C])	Low (65°F–75°F [18°C–24°C])
Processor Temperature	High (113°F–167°F [45°C–75°C])	Low (113°F–122°F [45°C–50°C])
Processor Temperature Fluctuation	High (113°F–167°F [45°C–75°C])	Low (113°F–122°F [45°C–50°C])

Source: Asetek, Inc.

Operational Testing

Task Objective

The goal of this subtask was performance-controlled load testing for the liquid-cooled racks operating under the optimum external-loop water settings, to provide a quantitative comparison of improved performance from liquid cooling versus air cooling.

The team was unable to perform controlled load testing at 25 percent, 50 percent, and 75 percent. However, the team was able to model the performance of clusters at these loads, and the results follow.

Benchmarking Study

kW Engineering built energy models to determine the energy consumption and efficiency associated with the Asetek systems installed at LLNL's B451 and B654 data centers. The models were calibrated based on power measurements and all trend data available. Asetek then modified these models to determine the energy consumption and efficiency at various cooling loads.

Asetek considered 100 percent load to be the maximum IT load measured for a sustained period during the trend period. Asetek then modeled the systems at 25 percent, 50 percent, 75 percent, and 100 percent load.

Site1: Building 451 at Lawrence Livermore National Laboratory

Table 15 summarizes the modeled B451 data center benchmarking results before and after the Asetek system was installed.

Table 15: B451 Benchmarking Summary

Load	Total IT Load (kW)	Baseline: Heat Removed by Asetek (kW)	Baseline: PUE	Post-retrofit: Heat Removed by Asetek (kW)	Post-retrofit: PUE
25 percent	237.6	-	1.66	35.7	1.65
50 percent	475.2	-	1.63	71.5	1.61
75 percent	712.8	-	1.66	107.2	1.63
100 percent	950.4	-	1.68	142.9	1.63

Source: Asetek, Inc.

Asetek assumed that, as the IT load increased, the distribution losses remained at 2 percent of the IT load, lighting energy stayed the same, and the quantity of constant speed at which the CHW pumps operated increased from one to three. Asetek assumed constant condenser water plant and chiller efficiency. Asetek assumed the fraction of heat removed by the Asetek system remained the same.

The PUE did not change much with load because most of the energy consumption was associated with the CHW and condenser water plants. Asetek assumed the change in load did not significantly affect the chiller or the condenser water plant.

CPU Performance

CPU and rack temperatures were evaluated for the Cab cluster in Building 451. Dr. Ghaleb and his group observed CPU and track temperature variations, but this was not done at varying CPU load. Drawing from experience, the processors would operate at lower temperatures as the load decreases, and the temperature fluctuations for the process would also decrease as load decreases. Asetek expected a corresponding increase in performance for the liquid cooling solution.

Site 2: Building 654 of Lawrence Livermore National Laboratory

Table 16 summarizes the modeled benchmarking results for the AHU system and the Asetek system. The AHU system includes just the AHU and exhaust fan power. The Asetek system includes the fluid cooler fan, pumps for the fluid cooler, and a pump for the CDU, all serving the Asetek system. The pPUE (partial PUE) is listed instead of PUE because it is limited to specific systems, not the entire data center.

Table 16: B654 Benchmarking Summary

Load	IT kW	AHU pPUE	Asetek pPUE
25 percent	346	1.04	1.56
50 percent	693	1.04	1.38
75 percent	1,039	1.08	1.33
100 percent	1,386	1.15	1.31

Source: Asetek, Inc.

Note that Asetek did not include the energy consumed by the rack pumps in the heating, ventilation, and air conditioning (HVAC) load because they were replacing the server fans and considered IT load.

As Table 16 shows, the pPUE for the Asetek system decreased as the load increased. This is because the two pumps operated at essentially constant flow. As the load increased, the associated pump energy stayed the same, decreasing the PUE.

The pPUE for the AHUs demonstrated an inverse relationship with the load because, as the load increased, the supply and exhaust fan speed increased. Due to the fan affinity law, a small increase in fan speed results in a large increase in power. This analysis assumes a minimum supply fan speed of 50 percent, based on experience.

CPU Performance

There is a correlation between processor temperature fluctuations and the average power leakage. Since the processor generated less heat and required less power at lower load, Asetek anticipated that leakage current would reduce accordingly. Building on the work done by Sarood et al. (2013), Asetek found that, for every 10°F (–12°C) increase in processor temperature, there was a doubling of the soft error rate. Asetek expected that the soft error rates would reduce at lower load, because the processor would operate at a lower temperature with less fluctuation.

Conclusion

The model built by KW engineering was used to predict the post-retrofit-controlled load and the production performance of the system.

Greenhouse Gas Reduction

Task Objectives

The goal of this task was to prepare a Lifecycle Greenhouse Gas (GHG) Analysis report that described the method used to determine GHG emission reductions for the project and the result. For the purpose of this report, Asetek focused on the GHG savings from the retrofit efforts on Building 451 because it provided the better comparison on the advantages of a liquid-cooled solution over an air-cooled solution. The installation in Building 654 was not an ideal site for comparison. The air-cooled Quartz cluster and the liquid-cooled Topaz cluster were installed with different IT loads, and the building cooling did not sufficiently delineate between the clusters, making a comparison challenging.

Energy Models

Asetek developed a GHG emission model for B451 using data collected from Task 7, as well as information from the CEC and the U.S. Environmental Investigation Agency (U.S. EIA). For the purpose of this report, Asetek considered only CO_2 emission reduction. Asetek performed the CO_2 emission analysis using two different scenarios:

- Scenario A: California CO_2 emission reduction
- Scenario B: United States CO_2 emission reduction

Data Collection

In addition to the data collection described earlier in this report, Asetek collected the following data to complete the CO_2 analysis:

- Average grams of CO_2 emissions per kWh consumed in California
- Average grams of CO_2 emissions per kWh consumed in the United States

Calculation Methodology

As explained earlier, the savings for the Asetek retrofit came primarily from reducing load on the CRAH units, allowing for their fans to ramp down the required fan energy. In addition, the reduced fan speed resulted in less heat being rejected to the CHW loop, thereby reducing the energy consumption by the CHW and condenser water plants.

Asetek determined the CO_2 emission reduction by calculating the total power demand of the data center; it assumed a data center lifetime of five years to compute the total energy consumed over the lifetime of the data center. Asetek then used the average CO_2 emissions per kWh in California to determine the CO_2 emissions from the data center pre- and post-retrofit. The difference between the two numbers was calculated to show the GHG reduction in California.

A similar approach was used to determine the CO_2 reduction using the national average. However, Asetek had to calculate the CO_2/kWh value by using U.S. EIA's value of total U.S. generation, and a Statista value of U.S. annual CO_2 emissions. The difference between the pre- and post-retrofit numbers was calculated to determine the CO_2 savings from the retrofit.

Result

From data collection and calculation, the Asetek system saves 355 megawatt hours (MWh) of electricity, or 3 percent of total data center electricity consumption, annually, which translates to an energy saving of 1,775 MWh of energy savings over the lifetime of a retrofitted data center with an IT load of 909.5 KW. Using California CO_2 emission numbers of 841 grams of CO_2 emissions per kWh, the Asetek system saves 298 metric tons of CO_2 annually and 1,490 metric tons of CO_2 over the lifetime of the data center. Using the U.S. CO_2 emission numbers of 1,342 grams of CO_2 emissions per kWh, the Asetek system saves 476 metric tons of CO_2 annually and 2,380 metric tons of CO_2 over the lifetime of the data center. This translates to an annual CO_2 emission reduction potential of 0.5 metric tons per KW in the U.S. and 0.3 metric tons per KW in California of IT load when retrofitted with the Asetek liquid cooling system.

CHAPTER 4: Technology/Knowledge/Market Transfer Activities

LLNL installed Asetek direct-cooling systems at two data centers on its campus, located at Building 451 (B451) and Building 654 (B654). The Cab cluster installed in B451 was a field retrofit to the existing Appro system. The liquid cooling equipment for the Topaz cluster was pre-installed at Penguin Computing (system manufacturer) before being delivered to LLNL. This chapter outlines the technology and knowledge transfer activities that occurred for each cluster.

Building 451

As LLNL has had extensive experience with liquid-cooled HPC clusters, Asetek primarily advised LLNL on specific facility requirements needed for Asetek equipment to perform properly. Further training on product operation was also provided as the CAB cluster was being retrofitted.

Topics included:

- Theory of Operation
- Asetek RackCDU™ Monitoring System Overview
- Service and Maintenance Activities on Asetek RackCDU™

Building 654

The Asetek liquid cooling equipment was installed at the Penguin Computing facility. Asetek provided guidance to Penguin Computing on the following topics:

- Installation Procedures — both cooling loops and RackCDU™
- Facility Requirements to Test
- Theory of Operation
- Asetek RackCDU™ Monitoring System Overview
- System Shipping Recommendations
- Service and Maintenance Activities on Asetek RackCDU™

Penguin Computing has implemented several other clusters with the same liquid-cooling architecture to other sites, using the knowledge gained from this installation.

Technology and Knowledge Transfer Activities

Knowledge gained from the project was made available to the public through a variety of sources during the grant period, including the Asetek website (https://www.asetek.com/videos/asetek-rackcdu-with-d2c-direct-to-chip-liquid-cooling/) conferences, presentations, and discussions with partners, utilities, and customers. Asetek has discussed the lessons learned with data center operators around the state and country in person, with the goal of implementing lessons learned in future installations.

The major groups influenced by this study are:

- Data center operators
- Original equipment manufacturers (OEMs)
- Investor-owned utilities (IOUs)
- The public

Trade Shows and Industry Conferences

Asetek is a regular participant in industry standard trade shows and has had exhibits annually throughout the grant period at:

- SC’12–SC’18
- ISC’13–SC’18

Asetek remains committed to a presence at trade shows and will continue to discuss the benefits of liquid cooling with a wide variety of customers and industry partners.

News and Press Releases

Asetek maintains a robust database of articles, publications, statistics, and other collateral on its website. Articles discussing the technology and its implementations can be found here:

- News — https://www.asetek.com/press-room/news/
- Company Blog — https://www.asetek.com/press-room/blog/
- In the Press — https://www.asetek.com/press-room/in-the-press/
- White Papers and Reports — https://www.asetek.com/press-room/media/white-papers-reports/

CHAPTER 5: Conclusions/Recommendations

Lessons Learned

In the B451 case, the heat was rejected to the same CHW 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 CHW pumps from the heat rejected by the Asetek system, significantly improving efficiency and dropping PUE. LBNL's study (Coles and Greenberg, 2014) found that, at 50 percent server load, it was possible to achieve a 14 percent overall site energy reduction for dry cooler rejection with chiller boost (with 104°F [40°C] supply water —minimizing water use for cooling and maximizing the potential for heat recovery) and a 20 percent overall site energy reduction for cooling tower heat rejection with chiller boost (with 68°F [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. By taking full advantage of the Asetek water cooling system, one can reduce or eliminate use of the chiller plant, significantly increasing the savings relative to that of the B451 case study.

In the B654 case, the Asetek system was compared to a compressor-free air-cooling system. The air-cooling system was highly efficient and operating in an appropriate range relative to design. In contrast, the water-cooling system operation was compromised both by over-sizing relative to the load studied and by its control scheme. In general, water cooling was significantly more efficient than air cooling, but the combination at B654 demonstrated that, without careful design and operation, it was possible for air cooling to be more efficient than water cooling. At B654, Asetek did not study the fact that air-cooled Quartz was throttled (reducing its useful computing) during high outdoor wet-bulb temperatures (which resulted in high supply air temperatures to the IT equipment), while Topaz was able to run unimpeded. Another advantage of the Asetek system that Asetek did not study was improvement to the computing efficiency. On the latter point, LLNL concluded that there was an approximately 7 percent computing efficiency improvement in the Cab cluster attributed to the Asetek system, due to cooler processor operation (Marathe et al., 2017).

In both the B451 and the B654 systems, an intermediate heat exchanger (the CDU) was used between the primary cooling water (that is, the CHW at B451 and the cooling water circulating through the fluid cooler at B654) and the cold side of the RackCDU™ (the hot side of which was where the heat was 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. They also impose an additional temperature difference, requiring the cooling water to be typically 2°F (-17°C) colder for the same temperature at the IT load; this decreases the efficiency of the heat rejection when compressor cooled and reduces the number of compressor-free hours (or forcing less-efficient operation) in systems with water-side economizers. Thus, eliminating the intermediate exchanger would save capital and operating costs, the latter in the form of energy savings. 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 the case of B654, eliminating the CDU exchanger and using the single smaller pump for cooling water would have reduced the energy for providing the cooling water by about 35 percent, which would have reduced the cooling energy overhead on the water cooling from about 30 percent to 19 percent.

The fluid cooler system at B654 operated at a relatively high approach to wet-bulb temperature, approximately 15°F (-9°C). A cooling tower paired with a plate-and-frame heat exchanger offers an approach of about half this much, which would have offered significantly better energy efficiency. On the other hand, the dry-cooling mode of the fluid cooler offers significant water savings when the ambient air temperature is low enough, leading to an energy/water tradeoff. These savings and tradeoff were not quantified as part of this study.

In short, the Asetek liquid cooling system and associated facility heat-rejection systems resulted in some energy savings at B451 (3 percent to 5 percent overall for the data center, and 7 percent for the Cab cluster) and negative savings at B654. In more typical applications, where advantage can be taken of the Asetek system to reduce or eliminate chiller operation, much higher savings could be anticipated, and positive savings would result from configurations like B654.

Summary and Conclusion

Building 451 System (Cabernet HPC Cluster Retrofit)

The engineering analysis that included high-level meter data for IT loads, mechanical loads including both the CRAHs and chillers, and other ancillary loads indicated 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 were equipped with a water–based Asetek cooling system.

From the statistical 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 1.64 to 1.59. However, when analyzing the performance of the Asetek cooling system that included only the CDU circulating pump and air-cooling systems for Cab, the pPUE was estimated to be 1.52, which would have been the expected PUE of the data center if all the servers had been cooled by Asetek. Asetek expected that this reduction would have been significantly larger if the Asetek system were not required to reject heat to the CHW system. The overall heat rejection fraction from Cab to the Asetek system was 37 percent, lower than in previous studies.

Building 654 System (Topaz HPC Cluster Compared to Quartz HPC Cluster)

The top-down analysis included the high-level meter data for IT loads and the mechanical loads that included both the water and the air-cooling systems. The resulting statistical model, while not fully complying with standards of validity, predicted an overall PUE of 1.15. The model also predicted a PUE of 1.17 for the hypothetical case of Topaz being 100 percent of the load and of 1.14 for Quartz being 100 percent of the load. Again, this model isn't fully valid, but it is consistent with higher efficiency for the air-cooling system than for the water-cooling one.

The bottom-up analysis, using a bin-data analysis to understand the performance of the Asetek system, the total overall data center, where Topaz and Quartz are the IT load, the overall PUE was 1.13. The pPUE for the air-cooled fraction of the load (most of Topaz and all of Quartz) was 1.07, and the pPUE for the water-cooled fraction of the load (the Asetek system and associated CDU, pumps, and fluid cooler) was 1.32. Air cooling was found to be more efficient than water cooling in this case. The overall heat rejection fraction from Topaz to the Asetek system was 23 percent, lower than in previous studies.

CHAPTER 6: Benefits to Ratepayers

Implementation of Asetek's proposed technology will benefit IOU electricity ratepayers in the near term by reducing the electricity bill for their data centers by up to 10 percent, based on the GHG reduction from power consumption reduction. As mentioned earlier, greater savings, on the order of 20 percent to 30 percent, may be possible if the chilled water system was not used for rejecting the heat from the Asetek system. Although Asetek did not perform peak load reduction analysis in this project, RackCDU™ could reduce load on the grid, reducing grid strain during peak times and improving grid reliability and safety, by cost effectively reducing statewide data center energy consumption. This project has the potential to cut statewide commercial electricity consumption by more than 2 percent if implemented across all data centers, reducing ratepayer electricity costs by lowering utility infrastructure expansion requirements and reducing the amount of electricity that must be imported from out of state.

Quantitative Benefits to California Investor-owned Utility Electricity Ratepayers

- Calculated Benefits (per data center): All data centers are different, and the actual energy saved will vary, based on location and data center design. In the GHG analysis previously mentioned, Asetek calculated Building 451 annual electricity savings as 354,780 kWh/yr and annual cost savings as $60K/year.
- Projected Benefits (statewide): RackCDU™ could reduce statewide data center annual energy consumption by as much as 30 percent, representing a total savings of 2,400 GWh/year of electricity savings, or 3.5 percent of California's total commercial electricity consumption. The statewide cost savings could be more than $340M per year, and the GHG savings could be more than 700,000 metric tons CO_2e/year.

Affected Market Segments in California (Size and Penetration Rates)

The team found that RackCDU™ is most cost effectively implemented as a pre-installed solution and demonstrates the greatest savings in high-use data centers running HPC and in high-density data centers. However, the key lesson from this project is that, while there is improvement in performance and reduction of power consumption, the retrofit process is too disruptive to the data center operation. Hence, deploying servers that have pre-installed liquid cooling loops will be most advantageous. Thus, while the total California market size for RackCDU™ is the stock of servers, existing and new for HPC data centers, penetration will likely be low outside of new servers.

Qualitative Benefits to California Investor-owned Utility Electricity Ratepayers

Reduced fan loads in servers and in data centers can reduce noise levels by almost 10 times, providing improved health and safety due to a reduction in hearing injuries, with the same penetration rates as previously indicated.

Cost-to-benefit Analysis for Project Costs

Total cost of ownership (TCO) was calculated by comparing the sum of energy cost savings, changes in operating costs, and the capital costs of installing the liquid cooling equipment with the costs of operating similar equipment with air cooling alone. The building energy models developed in Task 3 provided the energy cost savings.

The capital costs of liquid cooling include both the products provided by Asetek, the equipment necessary to transport the heat collected with the liquid at the racks to the point of heat rejection, and any special heat rejection equipment required to release the heat into the atmosphere. As the Asetek products are now commercially available, the prices for these products will be used as the capital cost. Note that the costs for this infrastructure are very much site-specific and, as such, projecting these costs directly into those that might be incurred in each different facility is risky. Excluded from these costs is equipment used specifically for measurements to conduct this study, which would not be included in standard operation.

Some of the capital equipment deployed has an expected life of 20 years. The Asetek equipment installed in the servers is unique to each server model and, as such, has the same life as the servers. The normal life of a server is five years. This server life is driven by hardware obsolescence rather than wear. Improvements in computing capability follow Moore's law, and after five years the new computers available are so improved that it is more efficient to replace the old with new servers. This refresh cycle assumes that the new servers are manufactured with liquid cooling.

The TCO model also captures any changes in non-energy operating costs that occur with liquid cooling, such as staff training and equipment maintenance.

The most straightforward way to evaluate TCO is with marginal analysis. Marginal analysis removes the cost of computer equipment from the calculation and focuses only on the cost of installing and operating the liquid cooling equipment, in contrast with the cost of operating air-cooling equipment and any additional air-cooling systems that may be required to support increased computing capacity.

B451 Cost-to-benefit Analysis

In this study, liquid cooling was installed in two data centers, one in Building 451 (B451) and one in Building 654 (B654). As stated earlier, the air-cooling strategies used in these two data centers were markedly different. B451 used traditional air cooling. Servers reject their heat into data center air, which is then cooled by CRAHs that capture the heat into chilled facilities liquid. Heat in the facilities liquid is then rejected by chillers into condenser water, and campus level cooling towers reject the heat from condenser water into the atmosphere. Energy is consumed in this process by the fans in CRAH units that circulate air to the servers, the pumps that circulate facilities liquid between CRAH units and the chillers, the chilling (refrigeration) process of the chillers, the pumps that circulate condenser water between the chillers and cooling towers, and the cooling towers.

The chillers for B451 had excess capacity. To minimize the cost of facilities equipment upgrades, Asetek rejected heat from the liquid cooling system into the facilities chilled water loop. This decision had two consequences.

1. A third-party CDU was required between the facilities chilled water loop and the "process" water loop that collected heat from the Asetek liquid cooling equipment, to ensure that the process water remained above the dew point in the data center.
2. It limited the energy-saving potential of the installation because the chillers (which were a major consumer of cooling energy) were still being used to reject the heat from the building into condenser water.

Greater energy savings would have been achieved if the direct-to-chip liquid cooling equipment had been connected directly to the condenser water loop. However, the capital cost of deploying the CDU and process water loop in B451 was quite low, totaling $113,938. Of this, $53,780 was for the purchase of the CDU that separated chilled water from process water to avoid condensation risk within the servers.

B654 Cost-to-benefit Analysis

B654 was built as a state-of-the-art air-cooled data center. Evaporatively cooled air is supplied to the data center by large outdoor air handling units (AHUs). No chilling of the air is required. These AHUs were built (reconditioned) primarily for this purpose, and the design of the units focused on their efficiency in performing this function. These units include the ability to evaporatively cool water. However, this ability was not "economized" and thus the evaporative cooler consumes fan energy at all times. Further, the facilities loop within B654 was "future proof" constructed to support much larger loads than were used in this study. But the design did not allow the system to be "turned down" to more efficiently cool the load used. This resulted in the liquid cooling system in B654 consuming more power than the air-cooling system. Given this circumstance, the TCO for B654 is negative, as there are no energy savings to offset the cost of liquid cooling. However, it is important to note that LLNL continues to use Asetek liquid-cooled computing in B654. This will lower the cost of liquid cooling within this data center. In general, Asetek expects liquid cooling to have lower operating costs than state-of-the-art air-cooled data centers. However, given the overall efficiency of both state-of-the-art air-cooled and liquid-cooled data centers and the extra cost of liquid cooling computers, it will likely be factors other than cooling energy cost that drive the choices between liquid and air cooling, most notably the ability to deploy high interconnect density systems with liquid cooling and the superior stability of the computer when liquid-cooled, as demonstrated in the Computational Performance Analysis section of Chapter 3. The cost to put the facilities liquid cooling infrastructure into B654 was $953,000. This excludes the existing AHU with evaporative water-cooling capability.

Project Costs

Turning to the numbers, the commercial cost of the direct-to-chip liquid cooling system used to cool Topaz was $268,654. This cost cooled 754 servers installed in 12 open compute project racks, making the per-server cost $356.31 and the cost per rack $22,388.

In B451, Asetek performed a retrofit on existing servers. Experience in this retrofit indicates that retrofitting existing servers with aftermarket liquid coolers is not an economically viable business model for liquid cooling vendors. There are a number of reasons for this. The installation cost in a retrofit is considerably higher than when liquid cooling is installed during server manufacture. In the retrofit model, the servers must be removed from the cluster, air-cooling heat sinks must be removed, the chassis modified, and liquid cooling installed; and then the servers must be returned to the rack. This process creates waste. The air-cooling heat sinks must be recycled, but recycling will not recover the original cost of these parts, the labor used to install them in the first place, or the labor required to remove them. Asetek also learned that some of the engineering checks and balances present when a liquid cooling solution is designed in cooperation with a server manufacturer are missing in the retrofit model. This can lead to design problems that appear only during the installation and test process of the retrofit. These problems delay the retrofit and cause additional rework. This is disruptive to data center operations. Thus, Asetek does not see value in using retrofits to gain energy efficiency moving forward. For this reason, rather than using retrofit costs in the TCO analysis, Asetek will apply the commercial costs observed with Topaz to the Cab cluster in B451. This will better represent the cost of future installations.

The Cab cluster consists of 1200 server nodes installed in 17 racks. The commercial price for the liquid cooler in each server node from Topaz is $110.67. Each server node needs a pair of connecting tubes to attach it to the RackCDU™. These cost $27.68 per server node. This gives a total of $134.35 per node, for a total of $161,220.

Each of the 1200 server nodes installed in the 17 racks must be equipped with a RackCDU™ and a pair of facilities hoses. The commercial price of the RackCDU™ from Topaz is $7,339.12 and the hose pairs are $920.00. Installation was $334.00 per rack. This installation includes training of LLNL staff on liquid cooling. This gives a total of $8,593.12 per rack, or a total of $146,083.04 for 17 racks.

Shipping and handling charges to get all the materials to LLNL amounted to $34,811.40 for 12 racks worth of gear, pro-rated across 17 racks for a total shipping cost of $49,316.15.

In sum, the total cost for the direct-to-chip cooling system to equip Cab with liquid cooling is $356,619.19. When the cost of facilities modifications of $113,938 is added, the initial capital cost totals $470,557.19.

On the cost recovery side, data centers have energy savings of 348,663 kWh per year. The U.S. EIA put the average cost of electricity to industrial customers at 11.23 cents per kWh. This results in energy cost savings of $39,154.85.

At these energy savings, the payback period is 12 years. Unfortunately, the useful life of a server is in the range of five to seven years. This will likely force at least one refresh of the servers, which will include replacing the liquid coolers within the servers, as these cannot be moved from one server to the newer-generation server. The server refresh will result in an additional capital expenditure of, at minimum, the cost of the liquid coolers in the server, or $132,804 (assuming the same number of servers is purchased in the refresh).

GLOSSARY AND LIST OF ACRONYMS

Term	Definition
AHU	air-handling unit
BAS	building automation system
BIOS	basic input/output system
Btu	British thermal unit
°C	degree(s) Celsius
Cab	Cabernet
CDU	cooling distribution unit
CEC	California Energy Commission
cfm	cubic feet per minute
CHW	CHW
CHWST	CHW supply temperature
CPU	central processing unit
CRAH	computer room air handler
CVRMSE	coefficient of variation RMSE
DIMM	dual inline memory module
DX	direct expansion
ECM	energy conservation measure
EVO	Efficiency Evaluation Organization
EPIC	Electric Program Investment Charge
°F	degree(s) Fahrenheit
GHG	greenhouse gas
gpm	gallons per minute
GWh	gigawatt-hour
hp	horsepower
HPC	high-performance computing
HVAC	heating ventilation and air conditioning
IOU	investor-owned utility
IPC	instructions per cycle
IPMI	intelligent platform management interface
IT	information technology
IWC	inches of water column

Term	Definition
kW	kilowatt
kWh	kilowatt-hour
LBNL	Lawrence Berkeley National Laboratory
LLNL	Lawrence Livermore National Laboratory
M&V	measurement and verification
MW	megawatt
NBE	net determination bias error
NREL	National Renewable Energy Laboratory
OAT	outside air temperature
OEM	original equipment manufacturers
PDU	power distribution unit
pPUE	partial power usage effectiveness
PUE	power usage effectiveness
PUE_{BL}	baseline power usage effectiveness
PUEPR	post-retrofit power usage effectiveness
PUE_{Q}	quartz-power usage effectiveness
PUE_{T}	topaz-power usage effectiveness
R_2	coefficient of determination
RAT	return-air temperature
RMSE	root mean squared error
ROI	return on investment
SAT	supply-air temperature
TB	transformer bank
TCO	total cost of ownership
TMY3	Typical Meteorological Year, 3rd collection
U.S. EIA	United States Environmental Investigation Agency
UPS	uninterruptible power supply
VFD	variable frequency drive
W	watt
W/sf	watts per square foot

References

Coles, Henry C., and Steve Greenberg. 2014 (March). "Direct Liquid Cooling for Electronic Equipment." Report 6641E. Lawrence Berkeley National Laboratory. Available at https://eta.lbl.gov/publications/direct-liquid-cooling-electronic

EVO (Efficiency Valuation Organization). 2012 (January). "International Performance Measurement and Verification Protocol: Concepts and options for determining energy and water savings." Volume I, EVO 10000-1:2012. Efficiency Valuation Organization. Available at https://www.eeperformance.org/uploads/8/6/5/0/8650231/ipmvp_volume_i__2012.pdf

EVO. 2016 (October). "Core Concepts: International Performance Measurement and Verification Protocol." EVO 10000-1. Efficiency Valuation Organization. Available at https://evo-world.org/images/denisdocuments/Core-Concept-2016.pdf

Hughes, Phil, Henry Coles, Steve Greenburg. 2014 (October). "Rack-Level Cooling Using Cold Plates." *Data Center Handbook*. Chapter 26, 479-493. Hwaiyu Geng, ed. Wiley and Sons, Hoboken, NJ. Available at http://dx.doi.org/10.1002/9781118937563.ch26

Marathe, Aniruddha, Ghaleb Abdulla, Barry Rountree, and Kathleen Shoga. 2017 (May). "Towards a Unified Monitoring Framework for Power, Performance and Thermal Metrics: A Case Study on the Evaluation of HPC Cooling Systems." LLNL-Pres-593023. Lawrence Livermore National Laboratory. Available at https://www2.cs.arizona.edu/~amarathe/papers/marathe_hppac17.pdf

Sarood, Osman, Esteban Meneses, and Laxmikant Kale. 2013 (November). "A 'Cool' Way of Improving the Reliability of HPC Machines." Available at http://charm.cs.illinois.edu/newPapers/13-25/paper.pdf

Sickinger, David, Otto Van Geet, and Chris Ravenscroft. 2014 (November). "Energy Performance Testing of Asetek's RackCDU™ System at NREL's High Performance Computing Data Center." NREL/TP-2C00-62905. National Renewable Energy Laboratory. Available at https://www.nrel.gov/docs/fy15osti/62905.pdf