

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

Performance Evaluation for Modular, Scalable Liquid-Rack Cooling Systems in Data Centers

Permalink

<https://escholarship.org/uc/item/5zc7h0sz>

Author

Xu, TengFang

Publication Date

2009-08-13

**Performance Evaluation for Modular, Scalable Liquid-Rack
Cooling Systems in Data Centers**

Final Report

To

The California Energy Commission

May 2009

Tengfang Xu

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

Disclaimer

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

Contents

EVALUATION FOR A MODULAR, SCALABLE LIQUID-RACK COOLING SYSTEM IN DATA CENTERS	2
1 PROBLEM STATEMENT	2
2 TECHNICAL OBJECTIVES	2
3 TECHNICAL INFORMATION ON THE CHARACTERISTICS OF COOLING SYSTEMS, SERVERS	3
3.1 MODULAR LIQUID-RACK COOLING EQUIPMENT	3
3.2 SERVER	5
3.3 SERVER POWER MANAGEMENT.....	5
3.4 EQUIPMENT LOCATION	6
4 MEASURED PARAMETERS.....	8
5 DATA ACQUISITION SYSTEM.....	8
6 TEST PROCEDURES AND OPERATING CONDITIONS.....	9
7 PERFORMANCE METRICS FOR MODULAR COOLING	10
8 SUMMARY OF FINDINGS AND CONCLUSIONS.....	12
9 ACKNOWLEDGEMENTS.....	16
10 REFERENCES.....	16

Evaluation for a Modular, Scalable Liquid-Rack Cooling System in Data Centers

1 Problem Statement

Scientific and enterprise data centers, IT equipment product development, and research data center laboratories typically require continuous cooling to control inlet air temperatures within recommended operating levels for the IT equipment. The consolidation and higher density aggregation of slim computing, storage and networking hardware has resulted in higher power density than what the raised-floor system design, coupled with commonly used computer rack air conditioning (CRAC) units, was originally conceived to handle.

Many existing data centers and newly constructed data centers adopt CRAC units, which inherently handle heat transfer within data centers via air as the heat transfer media. This results in energy performance of the ventilation and cooling systems being less than optimal. Understanding the current trends toward higher power density in IT computing, more and more IT equipment manufacturers are designing their equipment to operate in “conventional” data center environments, while considering provisions of alternative cooling solutions to either their equipment or supplemental cooling in rack or row systems.

In the meanwhile, the trend toward higher power density resulting from current and future generations of servers has created significant opportunities for precision cooling suppliers to engineer and manufacture packaged modular and scalable systems. The modular and scalable cooling systems aim at significantly improving efficiency while addressing the thermal challenges, improving reliability, and allowing for future needs and growth. Such pre-engineered and manufactured systems may be a significant improvement over current design; however, without an energy efficiency focus, their applications could also lead to even lower energy efficiencies in the overall data center infrastructure.

The overall goal of the project supported by California Energy Commission was to characterize four commercially available, modular cooling systems installed in a data center. Such modular cooling systems are all scalable localized units, and will be evaluated in terms of their operating energy efficiency in a real data center, respectively, as compared to the energy efficiency of traditional legacy data center cooling systems.

2 Technical objectives

The technical objective of this project was to evaluate the energy performance of one of the four commercially available modular cooling systems installed in a data center in Sun Microsystems, Inc. This report is the result of a test plan that was developed with the industrial participants’ input, including specific design and operating characteristics of the selected modular localized cooling solution provided by vendor 3.

The technical evaluation included monitoring and measurement of selected parameters, and establishing and calculating energy efficiency metrics for the selected cooling product, which is a modular, scalable liquid-rack cooling system in this study. The scope is to quantify energy performance of the modular cooling unit in operation as it corresponds to a combination of varied server loads and inlet air temperatures, under various chilled-water supply temperatures.

The information generated from this testing when combined with documented energy efficiency of the host data center's central chilled water cooling plant can be used to estimate potential energy savings from implementing modular cooling compared to conventional cooling in data centers.

3 Technical information on the characteristics of cooling systems, servers

The evaluation tests were performed in a data center space located in Santa Clara, California. The datacenter area is approximately 12,800 square feet, with a ceiling height of 13 ft 6 inches and no raised floor. All server racks and support equipment are installed directly on the slab floor. There were various types of servers, rack sizes and shapes from various vendors. The data center was specifically designed to support racks with any type of IT equipment.

Power, chilled water, and communication cables to the server racks were provided through overhead cable trays. 700 tons of cooling is provided to the space from the central chilled water plant for cooling the IT equipment. The chilled water is supplied by a 2,000-ton central chilled water plant.

Power for thirty-six 150kW PDUs was provided through a 480v AC bus way system. The 150kW power distribution units (PDUs) were located throughout the space. The PDUs transform the power from 480v AC to 208v AC for distribution to the server racks. The data center currently can support 190 watts/ft² of floor area. The original design load per rack footprint in the existing data center was 5 kW/rack with growth to 9 kW/rack.

3.1 Modular liquid-rack cooling equipment

The liquid-rack cooling system studied consisted of a single modular system, sized 78"x12"x48" attached to two equipment racks that measured 78"x24"x48". Figure 1 shows the modular liquid-rack cooling enclosure. The modular liquid-rack system was powered via a single 208V single-phase 20A electrical circuit. Where redundant electrical service is available, two such circuits are available and provided. Water connections were made at the rear of the unit via two externally threaded 1" connections.

This modular, scalable, liquid-rack cooling system is designed to provide rack cooling up to 35 kW cooling per enclosure depending upon the supplied chilled water temperature and desired server air inlet temperature set point. When dealing with two racks of servers, such a liquid-rack unit is designed to provide cooling capability for higher server loads than those commonly seen in actual data center to date, e.g., up to 17.5 kW per rack. Using a local controller, the system ensures that all components are operating at preset

levels (e.g., inlet air temperature) to support the removal of specific heat loads. Three main features comprise the control system:

- Three (3) independent temperature zones (top, middle, bottom) with each zone monitoring inlet and outlet air temperature, water inlet and outlet temperature, fan speed and water flow rate.
- Variable speed fans with four (4) speed set points. The fans have adjustable speeds for given heat loads.
- Cooling water flow control from 0 – 100% of maximum flow rate. As with the fans, cooling water flow rate is regulated to provide the appropriate heat removal capacity based on the heat load.

With control and monitoring of all parameters regulated by a local controller, an end user is able to establish a specific set point for server's inlet air temperature. By constantly monitoring the various temperature and flow parameters, the controller regulates water flow and fan speed to achieve a preset supply air temperature.

In order to test the cooling performance and energy efficiency of the modular liquid cooling system with varying supply water temperatures, an additional coolant distribution unit (CDU) was installed and connected to the system. The CDU consists of a water-to-water heat exchanger that provided isolation between the primary chilled water loop and secondary cool water loop that is connected to the modular system. The CDU contained its own pump and monitoring equipment to measure primary and secondary chilled water pressure, temperatures, and flow rate.

Airflow through the modular system is horizontal such that cool air is pushed to the front of the server equipment, drawn through the equipment and exhausted where it is drawn back into the modular cooling system and passed through the unit's air-to-water heat exchanger to cool the supply air to the desired temperature.




Figure 1 Modular liquid cooling-rack enclosure

3.2 Server

The IT or heat load in each server rack was provided by 38 standard Sun V20 1U servers, each with a size of 28"x17"x1.75" each the two server racks tested. The maximum load per server rack that was actually achieved in this study was 9.75kW for all temperature set points tested. Details of the servers provided in this study are publicly available [1].

Two racks of servers are equipped with one liquid-rack cooling unit in this study due to applicability. Given that such a liquid-rack unit is designed to provide cooling capability of up to 17.5 kW per rack, and the maximal available server load being no more than 10 kW per rack, we only tested the liquid-rack cooling unit's cooling performance for its partial load operation in this study.

3.3 Server power management

Using a commercially available software program as the "control program", the load within each rack was effectively controlled at desired levels by dynamically turning on and off servers and running the program [2] at various CPU loads to achieve the desired power consumption and resulting heat load to test the energy efficiency of the modular cooling system. Prior to the test, reference measurements on each type of server that was being used in the racks were performed to measure idle and loaded power consumption.

To achieve the desired partial or full power load level (kW/rack) to be tested, the number of servers needed to run at 100% load was calculated and the program turned the rest off. For example to achieve the desired heat load levels in the test plan, the number of servers running at full available server load per rack were as follows: 10 kW - 33 servers; 7.5 kW - 25 servers; 5 kW - 17 servers; 2.5 kW - 9 servers.

Although all of the servers used in the test configuration were the same, the initial reference measurements identified that they had significantly different power consumption due to different memory or computing configurations installed.

Therefore, to achieve and maintain the desired full or partial power load per rack during each test sequence, the monitoring system collected real time measurements of server power from the rack power strips and the program used this information to turn on or off additional servers as necessary to maintain the desired power load levels.

To monitor the inlet air temperature being delivered to the test racks by the modular cooling system, air temperature probes were installed at the top, middle and bottom of each rack. To improve the response time of these sensors, the power to the servers installed at these rack elevations were maintained on during each testing sequence. Prior to starting a specific modular cooling system test sequence, the total power consumption at each rack was verified against the readings of the power strip and adjusted as needed until the power consumption was stable.

3.4 Equipment location

The two server racks were positioned at the southwest corner of the data center in this study. The server racks together with the modular system and CDU were separated from the rest of the data center by an array of curtains surrounding them. The space within the curtain has a floor area of 110” by 178” and a height of 86”. Figure 2 shows the front- and rear-view of the equipment layout in this study.

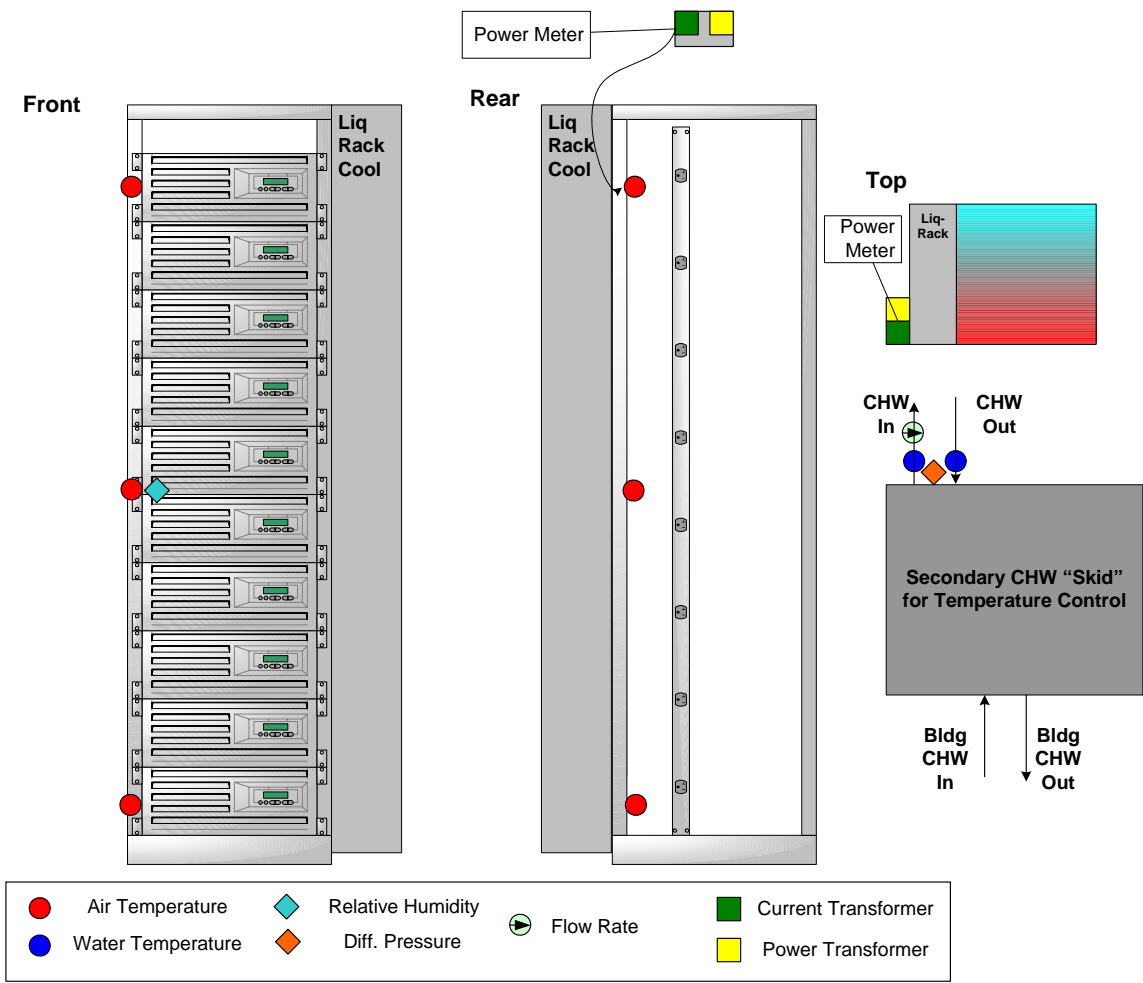


Figure 2 Liquid Rack Cooling Layout

4 Measured parameters

The following parameters were monitored or measured during the evaluation.

- Power demand of servers and cooling modules
 - Actual power demand for servers used in this study.
 - Actual power demand for the cooling modules
 - Electric power demand for the pump unit and eight cooling modules was monitored separately.
- Air temperature
 - Cold inlet air temperature to the front of server racks
 - There were three temperature sensors (RTDs) installed at the bottom, middle, and top positions (0.65”, 37.5”, and 69”, respectively) at the front inlet of each server rack. These heights corresponded to servers 2U, 20U, and 38U, respectively.
 - Hot outlet air temperature from the back of server racks
 - There were three temperature sensors (RTDs) installed at the bottom, middle, and top positions (0.65”, 37.5”, and 69”, respectively) at the rear outlet of for each of the server racks.
 - Data center air temperatures (outside the enclosed test area): from building energy management system
 - Outdoor air temperatures (dry-bulb & relative humidity).
- Cooling module entering and leaving chilled water temperatures
 - Chilled water temperatures in the supply and return pipes were measured by installing two temperature sensors on the surface of the water pipes, with insulation material wrapped around.
- Cooling module chilled water flow rates
- Cooling module entering and leaving chilled water pressure differential.

In addition to the real-time measurements taken of the test environment and the systems,, the following parameters were recorded manually to quantify the power demand in the data center and the energy use of the central chilled water plant: total power demand going into all IT equipment, and total power demand to chiller plant.

5 Data acquisition system

Data collection for the test environment was performed using a commercially available data acquisition system [3]. Data collection for the central chilled water plant was collected from the onsite energy management system or from field measurements . Data points were measured over the duration of the study, gathered from the manufacturer's modular cooling equipment (where available), the Sun servers, the rack power strips, and

an array of power meters, flow meters, pressure transducers, and RTD temperature sensors.

Data was gathered by local network appliances via a variety of network and serial communication protocols from the meters, the servers, and various analog sensors through I/O modules. After initial local processing and alarm checking, data was reported to a remote server and stored in a relational database. Similar data points were measured for each rack cooling technology, and stored in a shared relational database at a remote server. The real-time data was available through a web application, allowing users to monitor and manage the study remotely in real time. Access controls ensured that each manufacturer could see only its own data, while the designated host had access to all data.

Three power meters measured the energy use of the rack chiller and the two racks of servers. Smart power strips reported electric current for each rack. RTDs were placed at three heights on each rack, front and back, as well as at the inlet and outlet of the fan units. Ambient air temperatures and humidity were measured on the cold-aisle. Various internal server temperatures were gathered from selected servers, as reported by the servers themselves. Supply and return chilled water conditions were measured using a flow meter, pressure transducers, and RTDs. The cooling modules themselves also reported water conditions, fan speed, air temperatures, and cooling output power.

6 Test procedures and operating conditions

The following procedures were followed to control and maintain a pre-determined matrix of inlet air temperatures and server loads.

- Control and adjust the server's inlet air temperature at discrete set points from 80°F down to 68°F, with steps of 4°F.
- For each desired inlet air temperature level, first identify the available and applicable full server load (e.g., 9.75kW per rack). Then control the server load and allow sufficient operating time to perform steady-state testing and monitoring.
- Decrease the server load from full (100% available) to 50% of the defined capacity per rack (e.g., 9.75 kW/rack). Perform steady-state testing and monitoring for 50% load level.
- Vary chilled water supply temperature at three levels: 70°F, 60°F, and 50°F.

The matrix of various inlet air temperatures (from 68°F to 80°F) with various chilled water temperatures (from 50°F to 70°F) used for both full load and half load in this study is presented in the following table.

Table 1 Set Points for Test Conditions

Supply water temperature (F)	Inlet air temperature set point (F)	Targeted total server load (kW) for one servers rack	Various server loads (kW) per rack	
			100%	50%
50	68	10	9.8	4.9
50	72	10	9.8	4.9
50	76	10	9.8	4.9
50	80	10	9.8	4.9
60	68	10	9.8	4.9
60	72	10	9.8	4.9
60	76	10	9.8	4.9
60	80	10	9.8	4.9
70	-	10	9.8	4.9
70	72	10	9.8	4.9
70	76	10	9.8	4.9
70	80	10	9.8	4.9

7 Performance metrics for modular cooling

In order to characterize the thermal performance of the modular liquid cooling system, we calculated the ratio of cooling produced by the modular system to the total power demand for the modular system, defined as “coefficient of performance (COP).” Normally COP of a modular system is the ratio of the heat removed by the modular system to the work supplied to the modular system.

The COP is unit less, with a higher value representing higher efficiency for the modular system. The COP can be calculated under applicable operating and environmental conditions (a range determined by server power load and inlet air temperature).

In this evaluation, the work supplied is fan power required to produce the required airflow from the modular system, while the heat removed is equivalent to the cooling transferred by the modular system.

$$COP = \frac{Cooling}{P_{total}}$$

Where

Cooling is the cooling provided by the modular system, and

P_{total} is generally defined as the total power demand for all components (e.g., fan, and pump if necessary) in the modular system.

Because there are only fans in this modular system, the total power demand was for air-circulating requirement in this case.

Total power demand for the modular system can be measured as follows:

$$P_{total} = P_{fan\ power}$$

The actual cooling provided by the water-modular system can be calculated from the secondary-loop chilled water temperature rise and chilled water flow rate, using the following formula:

$$Cooling = \frac{60\rho Q C_p \Delta T_w}{3412.1}$$

Where

Cooling is produced from the modular system, in kW.

ρ : Water density in lb/gal, assuming water density ρ of 8.34 lbm/gal (or 62.4 lbm/ft³)

Q: Averaged water flow rate measured in gallon per minute

C_p: Specific thermal conductivity of water, 1 BTU/F-lbm

ΔT_w : Measured water temperatures rise, in °F

Therefore,

$$Cooling = 0.1467Q \Delta T_w$$

$$COP = \frac{Cooling}{P_{total}} = \frac{0.1467Q \Delta T_w}{P_{total}}$$

The portion of chiller pumping power required to deliver the chilled water volume in the primary-loop was ignored for this evaluation.

Another performance metric we calculated is the ratio of total power divided by the cooling transported by the modular liquid-rack cooling system, defined as modular system efficiency (MSE). This is similar to chiller efficiency commonly defined as power demand per cooling transferred. Represented in kW per cooling ton, a lower value of this ratio indicates a higher cooling energy efficiency at which the modular system is performing.

$$MSE = \frac{12000P_{total}}{Cooling} = \frac{12000P_{total}}{60\rho Q C_p \Delta T_w} = 24 \frac{P_{total}}{Q \Delta T_w}$$

Where:

Modular System Efficiency (MSE): ratio of total cooling power to cooling transported, in kW/ton

Q: Averaged water flow rate measured in gallon per minute

ΔT_w : Measured water temperatures rise, in °F

An alternative metric, defined as the modular system's power utilization index, is the ratio of power demand for the modular system to computer load under selected operating conditions. A higher value of the power index indicates higher cooling energy demand for the modular system at a given server load under an operation condition.

$$PI = \frac{P_{total}}{P_{server}}$$

8 Summary of findings and conclusions

The measurement and data collection system deployed in this study was reliable and accurate, and provided continuous monitoring of a wide range of critical parameters. It also provided real-time data display during the course of the experimental study. Data analysis was further enabled by writing custom database queries to parse the raw data collected to provide the ability for effective analysis of the large amount of data collected during the testing.

The software program used in the study to measure and monitor the power to the test environment effectively created various load/power consumption scenarios (based on the reference measures) to make sure the necessary power draw was generated and maintained required for all the tests in this study. Each rack was capable of providing approximately 10 kW and depending on the server load set points; the program was used to set load levels (e.g., 2.5kW, 5kW, 7.5kW, 10kW or maximum) by turning on/off the necessary amount of servers. It is noted that such a liquid-rack unit is designed to provide cooling capability of up to 17.5 kW per rack, with the maximal available server load being no more than 10 kW per rack, we only tested the liquid-rack cooling unit's cooling performance for its partial load operation in this study.

The overall coverage of operating conditions ranged from 67°F to 83°F for inlet air temperatures, the server loads ranging from 4.9 kW/rack up to 9.8 kW/rack, corresponding to elevated chilled water supply temperature ranging from 50°F, 60°F to 70°F.

The inlet air temperature distribution in this study exhibited significant uniformity for most of the test conditions, from the bottom to the top of the server racks, and from rack to rack. For example, the difference between maximum and minimum average inlet temperature ranged from less than 0.1°F up to 2°F, with the standard deviation ranging

from less than 0.1°F to 1.5°F. The temperature non-uniformity between the two racks tested was not noticeable because the temperature control appeared to be effective within its desired set point as shown in Table 1. Using a local controller, the system ensures that all components are operating at preset levels to support the removal of specific heat loads. It was also noted that the time it took to reach stable readings of air temperatures was very short, e.g., often within 10 minutes or less, after a set point was changed during the test. Compared to other modular cooling units, results from this evaluation show that it exhibited best ability) to control air temperature within the targeted range, i.e., significantly less time required for stable readings, and smallest deviations from the pre-defined set points during operation.

Table 2 shows the results from the tests performed at the facility. It includes measured parameters such as server load; average inlet/outlet air temperatures monitored at three different heights for both server racks; chilled water temperatures; cooling delivered by the modular system; power demand of the modular system; and three performance metrics in this study: COP; modular system efficiency; and PI.

Table 2 Actual Test Conditions and Results

Test Point	Server Load per Rack (kW/rack)	Average Inlet Air Temperature (F)	Average Outlet Air Temperature (F)	Chilled Water Temperature (F)	Chilled Water Flowrate (gpm)	Total Cooling (kW)	Total Power for Cooling Module (kW)	Total Power for Server (kW)	COP	ModEff (kW/ton)	Power Index
1	9.8	68	99	50	12.7	24.2	2.0	19.5	12.1	0.29	0.10
2	9.8	72	102	50	14.3	20.9	2.0	19.5	10.6	0.33	0.10
3	9.8	77	106	50	12.9	17.1	2.0	19.5	8.6	0.41	0.10
4	9.8	78	110	50	8.2	21.6	2.0	19.5	10.9	0.32	0.10
5	9.8	72	102	60	14.3	20.9	2.0	19.5	10.6	0.33	0.10
6	9.8	75	106	60	16.1	26.0	2.0	19.5	13.0	0.27	0.10
7	9.8	75	106	60	16.1	21.3	2.0	19.5	10.6	0.33	0.10
8	9.8	79	110	60	10.6	20.2	2.0	19.5	10.1	0.35	0.10
9	9.8	83	115	70	16.1	26.0	2.0	19.5	13.1	0.27	0.10
10	9.8	83	115	70	16.1	21.3	2.0	19.5	10.7	0.33	0.10
11	9.8	83	115	70	16.1	21.3	2.0	19.5	10.7	0.33	0.10
12	4.9	67	92	50	6.6	12.6	1.5	9.8	8.3	0.42	0.16
13	4.9	72	94	50	5.3	11.6	1.5	9.8	7.8	0.45	0.15
14	4.9	76	98	50	4.2	11.2	1.5	9.8	7.5	0.47	0.15
15	4.9	80	103	50	3.4	12.1	1.5	9.8	8.1	0.44	0.15
16	4.9	67	92	60	16.1	14.2	1.5	9.8	9.4	0.37	0.15
17	4.9	70	97	60	7.7	10.1	1.5	9.8	6.8	0.52	0.15
18	4.9	74	100	60	5.3	10.1	1.5	9.8	6.7	0.52	0.15
19	4.9	79	102	60	5.8	11.9	1.5	9.8	8.0	0.44	0.15
20	4.9	77	98	70	16.1	9.5	2.0	9.8	4.7	0.74	0.20
21	4.9	76	99	70	16.1	14.2	1.5	9.8	9.4	0.38	0.16
22	4.9	76	99	70	16.1	9.5	1.1	9.8	8.3	0.42	0.12

It is clear that different IT equipment operation and environmental operating conditions affected the cooling delivery efficiency of the modular cooling units. Specifically, variations in server power load and inlet air temperature have resulted in different COP, module system efficiency (kW/ton), and power index for modular Cooling System 3.

Generally, total modular cooling power demand stabilized at each of the server loads, e.g., mostly around 2.0 kW for the achievable full server load (i.e., 9.8 kW/rack), and 1.1 - 1.5 kW for half of the achievable server load (i.e., 4.9 kW/rack) in this study. The cooling power demand decreased when lower server loads were in operation.

Under the similar rack power load and the same chilled water supply temperature, the COP of the modular system tended to decrease with an increase in inlet supply air temperature, while modular system cooling kW/ton value exhibited a reversed trend. In addition, the PI values appeared to be virtually unchanged for a similar rack power load with various inlet air temperatures.

Under the different chilled water supply temperatures (50°F, 60°F and 70°F) and with similar operating and environmental conditions (rack power loads and inlet air temperatures), the performance metrics – COP, kW/Ton, and PI, remained largely unchanged. This was a result that power demand of the modular system remained largely the same under the same server load, regardless of variations in inlet supply temperature or chilled water temperature. Because the liquid-rack modular system was designed to handle cooling load of 17.5 kW/rack, the modular system appeared to easily provide sufficient cooling needed (up to 9.8 kW/rack) at various supply water temperatures in this study.

Overall, the MSE values (kW/ton) of Cooling System 3 ranged from less than 0.3 up to 0.74 kW/ton. They tended to increase with the decrease in rack power loads. With chilled water temperatures ranging from 50°F to 70°F, the COP values of the modular cooling decreased with the decrease in server loads – indicating a higher energy efficiency in transporting cooling when challenged with higher rack power loads in this study. The PI values did not exhibit clear correlation with the change in server loads.

It is clear that the liquid-rack modular liquid cooling system used in this study would have performed differently (e.g., more efficiently) if the available server loads were to be higher (e.g., up to twice higher). Nonetheless, it provided sufficient cooling to the server racks in all test conditions, while controlling inlet air temperature with significant thermal uniformity. With control and monitoring of all parameters regulated by a local controller, we were able to establish a specific set point for the desired server air inlet temperature, normally within a fairly short time (e.g., 10 minutes). By constantly monitoring the various temperature and flow parameters, the modular system's controller effectively regulated water flow and fan speed to achieve preset rack inlet supply air temperatures. The liquid rack cooling system also demonstrated superior capability and effectiveness in controlling inlet air temperature compared to the other systems tested in this study.

Overall, the liquid-rack cooling system was tested when operating at its partial capacity. The maximal server load available in this study was 9.8 kW/rack. To cool the servers operating with their maximal or partial heat density, the modular cooling system's COP values ranged from 5 up to 13; MSE values (kW/ton) from 0.27 up to 0.74 kW/ton; and PI from 0.1 to 0.2 - all exhibiting an energy efficiency level at par with traditional

CRAHs under the selected operating and environmental conditions (rack power load and inlet air temperature). As noted earlier, in this study we only tested the liquid-rack cooling unit's cooling performance for its partial load operation, which approximated half of its designed cooling capacity. Therefore, this set of measured efficiency values was penalized against this particular modular cooling because it is designed to handle much higher-server loads than most of the other modular cooling systems are.

It is thus noted that it would be premature to directly compare the values of performance metrics between this modular cooling system with the other modular, scalable cooling systems, because design intent and applicability are different. If this modular liquid cooling system had been tested with higher rack power loads (e.g., up to 20 kW/rack), the efficiency of the existing unit would have been increasing by a factor of 100% compared to the results from this study. It is clear that this liquid-rack cooling system is well positioned for effectively cooling higher-density servers in today's and future data centers.

Power for the variable speed fans in this system is controlled at four discrete levels corresponding to four speed set points. In order to improve MSE (kW/Ton) and COP values of the system, this study recommends that the vendor consider using electric commutation fans to provide continuous 0-100% adjustability in future generations of the product, allowing significant fan energy savings even at its partial load cooling.

The use of modular, scalable cooling systems, such as the unit in this evaluation, may increase the overall data center efficiency while maintaining effective control over inlet air temperatures at desired levels especially when optimized with the central cooling plant. The results of our testing support the hypothesis that this type of liquid cooling systems may be able to perform effectively without chillers using only evaporative cooling in cool, dry climates.

This evaluation does not include the assessment of the potential energy savings possible if this cooling system technology was used for the entire data center. The findings from this study indicate that by implementing rack-liquid cooling in lieu of traditional CRAH units, the overall kW/Ton in the data center with high-density server loads (e.g., over 15 kW/rack) could be reduced. In addition, this type of modular cooling system also provides increased flexibility in data center configuration and layout. Therefore, integration of such modular, scalable cooling systems within the "traditional" data center infrastructures should be to be carefully planned for successful and efficient cooling in data centers.

The overall energy demand for cooling server racks in a data center is largely affected by the efficiency of the central cooling system, such as chilled water plant or cooling tower plant. This study has demonstrated that the effective operation of rack-liquid cooling system did not necessarily require low-temperature water, such as what is typically from chilled water plants (e.g., 45°F). Although the overall energy demand will be also be affected, to some extent, by the modular cooling units installed within the data center, operating a chilled water system with elevated water temperature (e.g., cooling tower system, or chilled water system) in tandem with liquid-rack cooling units could significantly reduce energy required for data center cooling compared to operating a chilled water system with lower water temperatures (e.g., 45°F).

In summary, this study recommends that the reader consider not only the energy efficiency performance of the modular scalable cooling system, but also the system's design capability, its effectiveness to control and maintain server inlet air temperature (e.g., within ASHRAE recommended levels), and its potential dependence on other cooling or humidification in data centers. In order to quantify or estimate the impact of modular, scalable cooling systems on overall data center energy efficiency, one must also assess their integration with the rest of the data center eco-system, the temperature range of chilled water available from the plant (cooling tower or chiller), the local weather conditions, and the power density characteristics of the data center.

9 Acknowledgements

The author wishes to acknowledge Bill Tschudi of LBNL; Dean Nelson and his team from Sun Microsystems, including Brian Day and Mike Ryan; Ray Pfeifer of the Silicon Valley Leadership Group; Earl Sacerdoti and his team at Modius, Inc.; Clemens Pfeiffer and his team at Power Assure, Inc., and a number of anonymous industry participants for their assistance and review of this project report. The project was funded by the Industrial Section of the Public Interest Energy Research (PIER) Program of the California Energy Commission. This work was also supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

10 References

1. V20z Servers provided by Sun Microsystems. <http://www.sun.com/servers/entry/v20z/index.jsp>
2. SPEC (Standard Performance Evaluation Corporation) Power, Power Assure, Inc.
3. Modius OpenData® Data Center infrastructure Manager, Modius Inc.