

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

Performance Evaluation for a Modular, Scalable Passive Cooling System in Data Centers

Permalink

<https://escholarship.org/uc/item/1x3806nd>

Author

Xu, TengFang

Publication Date

2009-08-13

**Performance Evaluation for a Modular, Scalable Passive Cooling
System in Data Centers**

Final Report

To

The California Energy Commission

May 2009

Tengfang Xu

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

This work was supported by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

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

PERFORMANCE EVALUATION FOR A MODULAR, SCALABLE PASSIVE 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 PASSIVE MODULAR COOLING SYSTEM	3
3.2 SERVER	4
3.3 SERVER POWER MANAGEMENT.....	5
4 MEASURED PARAMETERS.....	5
5 DATA ACQUISITION SYSTEM.....	6
6 TEST PROCEDURES AND OPERATING CONDITIONS.....	6
7 PERFORMANCE METRICS FOR THE PASSIVE MODULAR COOLING	8
8 SUMMARY OF FINDINGS AND CONCLUSIONS.....	10
9 ACKNOWLEDGEMENTS.....	13
10 REFERENCES.....	13

Performance Evaluation for a Modular, Scalable Passive 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 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 passive, modular localized cooling solution provided by vendor 4.

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 passive, modular, scalable liquid cooling system in this study. The scope is to quantify energy performance of the modular cooling unit corresponding to various 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 Passive modular cooling system

This modular, scalable cooling system is essentially a heat exchanger utilizing water at a temperature higher than the room dew-point temperature to transport thermal loads. It is a passive cooling system delivering cooling from the heat exchanger. The passive cooling system studied consisted of a single, 77"H x 26"W x 5.5"D modular system (air-to-water heat exchanger) attached to rear door of the server rack. Since the rear-door heat exchanger replaces the standard rear door, the installed module adds 4" to the total rack depth, not 5.5". The modular system does not utilize electrical power to operate. Figure 1 shows the side view of the passive modular cooling system with rear door open.

The modular system has a water-cooled device that mounts on the rear of the server rack to cool the heated server exhaust air. A supply hose delivers *above room dew point temperature* (ADP) conditioned water to the heat exchanger. Another hose delivers warmed return water back to the water pump or chiller. In this report, we refer to this as a secondary cooling loop. The primary cooling loop supplies building chilled water to secondary ADP cooling loops and air conditioning units. The passive cooling module is designed to remove up to 102,400 BTU (30kW) of heat.



Figure 1 Side View of the Passive Modular system with Rear Door Open

In order to test the effectiveness of the modular system under various supply water temperatures, an additional coolant distribution unit (CDU) was installed and connected to the modular system. The CDU consists of a water-to-water heat exchanger that provides isolation between the primary chilled water and secondary cooling water loops that are connected to the module. The CDU utilizes an internal secondary loop pump and measures primary and secondary water pressures, temperatures and flow rates.

The passive rear-door modular cooling system was tested in a hot/cold aisle configuration. Ambient room air was drawn from the front of the servers by servers' internal fans. Then the air passed through the servers while being heated by the servers and finally exhausted to the rear of the rack.

3.2 Server

The IT or heat load in each server rack was provided by 40 standard Sun V20 1U servers, each with a size of 28"x17"x1.75" for the one server rack tested. The maximum load per server rack achieved in this study was 10.9kW per rack for all temperature set points tested. Details of the servers provided in this study are publicly available [1].

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 rack 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 beforehand and the control program turned the rest off. For example, to achieve the desired heat load levels in the test plan, the numbers of servers running at full 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 temperature of inlet air being delivered to the test rack by the passive modular cooling system, air temperature probes were installed at the top, middle and bottom of the 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 of the rack was verified against the readings of the power strip and adjusted as needed until the power consumption was stable.

4 Measured parameters

The following parameters were monitored or measured during the evaluation.

- Power demand of servers and modular system
 - Actual power demand for servers used in this study
 - Actual power demand for the CDU
- Air temperature
 - Cold inlet air temperature to the server rack
 - There were three temperature sensors (RTDs) installed at the top, middle, and bottom positions (0.65”, 37.5”, and 69”, respectively)
 - Hot outlet air temperature from the server racks
 - There were three temperature sensors (RTDs) installed at the top, middle, and bottom positions (0.65”, 37.5”, and 69”, respectively)
- 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, 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 both the on-site energy management system and from field measurements taken where continuous monitoring data points were not available. 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.

Two power meters measured the energy use of the cooling unit and rack 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 parametric test matrix with varying inlet air temperature (68°F, 72°F, 76°F, 80°F) against elevated supply water temperature (50°F, 60°F, 70°F) and achievable server loads (pre-determined with a maximum of 10.9 kW/rack) 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	10.9	5.5
50	72	10.9	10.9	5.5
50	76	10.9	10.9	5.5
50	80	10.9	10.9	5.5
60	68	10.9	10.9	5.5
60	72	10.9	10.9	5.5
60	76	10.9	10.9	5.5
60	80	10.9	10.9	5.5
70	-	10.9	10.9	5.5
70	72	10.9	10.9	5.5
70	76	10.9	10.9	5.5
70	80	10.9	10.9	5.5

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

- Control and adjust the cold aisle air temperature (server's inlet air temperature) at discrete set point of from 80°F down to 68°F, with steps of 4°F.
- For each desired inlet air temperature level, identify the level of available and applicable full server load (e.g., 10.9 kW for the rack tested). Allow sufficient operating time, and perform steady-state testing and monitoring.
- Decrease the server power from 100% (~10.9kW) to 50% (~ 5.5 kW/rack). Perform steady-state testing and monitoring at 50% load.
- Control secondary loop water supply temperature to 70°F, 60°F & 50°F.

Note that for the majority of the testing, we used space heaters to raise the temperature of air entering the inlet of servers. We observed that a reasonably small movement or change of the heaters could cause a noticeable change in the sensor temperatures. The temperature for the "entire" air volume was not uniformly adjusted however, which could lead to an inaccuracy of the "true" profile for actual room air temperatures. This setup could either overestimate or understate the performance of the water system since the system would have to remove more or less heat. Since we did not measure the system airflow and the added fan heat for each test condition, we could not reliably quantify the magnitude of the error or bias. Therefore, for the purposes of this evaluation we focused our analysis on waterside and power monitoring readings rather than airside calculations.

7 Performance metrics for the passive modular cooling

There is no active electric power demand for operating the passive heat exchanger device to cool servers in the rack. However, we regard the hydraulic power of the water flow within the secondary loop (connecting the rear door heat exchanger with the CDU) as the power demand needed for operating the modular cooling device in this evaluation.

In order to characterize its own thermal performance, we use the ratio of cooling delivered by the passive modular cooling system to the total hydraulic power demand for transporting the water flow in the secondary loop, defined as “coefficient of performance (COP).”

Normally COP of a modular system is the ratio of the heat removed by the module to the work supplied to the module. The COP can be calculated under applicable operating conditions (a range determined by inlet air temperature, rack server load, and elevated supply water temperature).

In this evaluation, the work supplied is narrowly defined as the hydraulic power required for transporting the water flow rate through the passive modular cooling system (assuming 100% transformed energy efficiency), while the heat removed is equivalent to the cooling provided by the rear door. Hydraulic COP can be quantified by the following equation.

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

Where

Cooling is the cooling provided by the passive modular cooling system; and

$P_{hydraulic}$ is calculated by the multiplication of the water differential pressure flow rate using the following equation:

$$P_{hydraulic\ power} = 0.43506 Q\Delta P/1000, \text{ kW}$$

Where

ΔP : pressure differential between secondary-loop supply and return water flow, in psi.

Q: Averaged water flow rate measured in gallons per minute

Because the CDU was only used to provide variable varied chilled water supply temperature (e.g., 50°F, 60°F, or 70°F), the actual power required for the CDU operation is not considered to be part of power demand for the passive modular cooling system evaluated in this study.

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

Where

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

Where

Cooling is delivered by the passive modular cooling 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,

$$Cooling = 0.1467Q \Delta T_w$$

$$COP = \frac{Cooling}{P_{total}} = 337.1 \frac{\Delta T_w}{\Delta P}$$

Because this passive modular cooling system requires no internal fans or additional pumping energy, its COP is proportional to the ratio of the secondary water temperature differential to the secondary water pressure differential. The tests confirmed the hypothesis for a sensible cooling coil such as the passive modular cooling system provided by vendor 4. The chiller pumping power required to deliver the chilled water volume in the primary-loop was ignored for this evaluation in keeping with the evaluation of the other vendor's systems.

Another performance metric is the ratio of total power divided by the cooling transported by the passive modular cooling system, defined as modular system efficiency (MSE). This is similar to chiller efficiency 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.

$$MSE = \frac{12000(0.43506 Q\Delta P)/1000}{Cooling}$$

$$MSE = \frac{12(0.43506 Q\Delta P)}{60\rho QC_p \Delta T_w} = 1.0431 \times 10^{-4} \frac{\Delta P}{\Delta T_w}$$

where

Module System Efficiency (MSE): Ratio of total hydraulic power for cooling to cooling delivered, in kW/ton

Q: Averaged water flow rate measured in gallon per minute

ΔP : Measured water pressure difference between secondary-loop supply and return, in psi. Cooling is delivered by the passive modular cooling system, in BTU/hr.

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

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

8 Summary of findings and conclusions

The facility measurement system deployed in this study was reliable and accurate, and provided wide range of critical parameters under monitoring. It has also provided real-time data display during the course of the experimental study, and supported the study's data analysis effectively.

The software program used in the study 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 consuming over 10 kW and depending on the server load set points; the program was used to set load levels by turning on/off the necessary amount of servers and by running the benchmarking tool at full load.

Based on the actual sever load achieved, the level of inlet air temperature set points, actual water supply temperatures, and the available temperature control within the test space, the actual operating conditions may and may not be fully achieved as planned in the set point matrix shown in Table 1.

The actual operating conditions achieved in this evaluation are shown in Table 2. The actual inlet temperature was not able to go down to as low as 68°F as originally planned based on the 'whole room' temperature of approximately 70°F. The overall coverage of operating conditions ranged from 70°F to 83°F for inlet air temperatures, from 4.8 kW/rack to 11.2 kW/rack for server loads, and from 50°F, to 60°F and 70°F for supply water temperatures.

Table 2 shows the results from the tests performed at the facility, which include rack power load; average inlet/outlet air temperatures monitored at three different heights for the server rack; chilled water temperatures; cooling delivered by the passive modular cooling system (heat exchanger); hydraulic power demand of the modular system; and three performance metrics: hydraulic based COP; module system efficiency; and power index.

Table 2 Actual Test Conditions and Results

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)	Hydraulic Power for Cooling Module (kW)	Total Power for Server (kW)	Hydraulic power based COP	Hydraulic power based ModEff (kW/ton)	Hydraulic power based Power Index
108	70	92	52	99	81	0.047	108	171.7	0.02	0.004
108	73	120	51	102	93	0.049	108	1901	0.02	0.005
109	76	121	53	97	91	0.046	109	1962	0.02	0.004
112	82	129	51	11.1	122	0.053	112	2299	0.02	0.005
108	70	117	60	91	63	0.043	108	1441	0.02	0.004
108	74	120	60	91	69	0.043	108	1594	0.02	0.004
109	76	122	60	90	74	0.039	109	1888	0.02	0.004
111	81	129	60	92	87	0.040	111	2192	0.02	0.004
109	70	117	70	97	51	0.042	109	1214	0.03	0.004
109	73	120	70	97	57	0.046	109	1226	0.03	0.004
111	82	128	70	89	69	0.043	111	1625	0.02	0.004
112	82	129	60	92	89	0.044	112	2023	0.02	0.004
57	73	83	49	107	60	0.051	57	1165	0.03	0.009
58	73	118	50	94	51	0.045	58	1134	0.03	0.008
58	79	121	50	98	56	0.047	58	1196	0.03	0.008
58	83	122	49	11.2	72	0.054	58	1349	0.03	0.009
58	73	116	60	90	40	0.043	58	920	0.04	0.007
58	73	116	60	89	39	0.043	58	920	0.04	0.007
58	78	121	60	89	42	0.043	58	981	0.04	0.007
48	83	121	60	90	41	0.043	48	950	0.04	0.009
58	73	116	70	94	29	0.045	58	644	0.05	0.008
58	73	116	70	95	34	0.046	58	736	0.05	0.008
58	79	121	70	93	35	0.044	58	797	0.04	0.008
66	83	122	70	94	34	0.045	66	766	0.05	0.007

It is clear that different IT equipment and environmental operating conditions affected the cooling delivery efficiency of the modular cooling unit. Specifically, variations in server power load and inlet air temperature have resulted in different hydraulic COP, modular system efficiency (kW/ton), and power index for passive modular cooling system in this study, as exhibited for each of the chilled water supply with different temperatures.

Generally, the achievable full server load was 11.2 kW/rack, with partial load being as low as 4.8 kW/rack in this study. Total cooling hydraulic power remained very low and stable at each of the rack power loads; varying between 0.04-0.05 kW for full rack power load and half load, respectively. Under a similar rack power load and the same chilled water supply temperature, the hydraulic COP of the passive modular cooling system tended to increase with an increase of inlet air temperature, while the module's cooling MSC values (kW/ton) remained low and stable. Similarly, the PI values were also very

low and remained virtually unchanged for a similar rack power load with various inlet air temperatures. On the other hand, the total cooling delivered by the module tended to increase when inlet air temperature rose from 70°F to 80°F under similar server loads.

In this study, we have examined how this passive modular cooling system performed under different chilled water supply temperatures. Specifically, the performance of the passive modular cooling system with the range of 50°F-70°F supply water was tested and evaluated for full and partial rack power loads, by the installation of a secondary water loop. Under different chilled water supply temperatures (50°F, 60°F, and 70°F) and similar operating conditions (e.g., rack power loads); the hydraulic COP tended to decrease with an increase in inlet air temperature, while the MSC values (kW/ton) and PI values remained basically unchanged.

Hydraulic COP decreased as the rack power load decreased. In general, the MSC values (kW/ton) and PI values increased with an increase in rack power loads, with the magnitude of change in kW/ton of 0.02 to 0.05 kW/ton, and PI of 0.004 to 0.009. This indicated that rather insignificant power was required for transporting chilled water flow through the passive rear door heat exchanger system.

We also found that at elevated supply water temperature (50°F-70°F), the passive modular cooling system performed well, delivering reasonable cooling to the server rack with partial server load as well as with full server load in this study. The testing indicates that this technology is significantly more energy-efficient than traditional CRAH cooling units.

Overall, the hydraulic COP ranged from 64 to 220; module system efficiency (kW/ton) values ranged from 0.02 to 0.05 kW/ton; and PI was constant at less than 0.01, exhibiting highly energy efficient operation under the various operating and environmental conditions (rack power load, inlet air temperature load and supply water temperature) tested in this study. The passive cooling system performed effectively with higher supply water temperature (from 50°F to 70°F) tested in this study. Because the power demand for this passive modular cooling system was estimated based upon hydraulic power, a direct comparison of this set of metrics with those of other modular cooling systems would not be necessarily appropriate. Instead, the purpose of performing such calculations for the energy metrics based upon the hypothetical hydraulic power in this study is to compare the passive system's performance under various operation conditions (inlet air temperature, server load, and chilled water temperatures) when applying this system into data centers.

It should be noted that this cooling technology relies on the ability of the server fans to support air movement and overcome the static pressure drop across the heat exchanger. In this study, the cooling ability of the passive cooling system was tested only for cooling requirements of the IT equipment housed in the selected rack, assuming the ambient air temperature within the data center being neutralized by other means. In fact, its cooling effectiveness in the entire data center would also be influenced by the surrounding conditioned air that might require additional cooling from other systems within the data center. Furthermore, it should be noted that the passive system in the study exhibited no direct capability or effectiveness in the control of inlet air temperatures due to its design. This may inherit potential risks of compromising the ability to control air temperatures

effectively within the entire data center, if the passive system is not implemented or integrated with the rest of the data center systems carefully.

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, while the use of modular, scalable cooling systems, such as the unit in this evaluation, may increase the overall data center efficiency especially when optimized with the central cooling plant. The results of our testing support the hypothesis that these types of liquid cooling systems may be able to perform effectively without chillers using only evaporative cooling in cool, dry climates.

The findings from this study indicate that by implementing liquid cooling in lieu of traditional CRAH units, the overall kW/Ton in the data center could be reduced. 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 and evaluated for effectively cooling in data centers.

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 is 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-AC02-5CH11231

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.