

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

Performance Evaluation for Modular, Scalable Overhead Cooling Systems In Data Centers

Permalink

<https://escholarship.org/uc/item/39b221zw>

Author

Xu, TengFang T.

Publication Date

2009-08-13

**Performance Evaluation for Modular, Scalable Overhead 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 MODULAR, SCALABLE OVERHEAD COOLING SYSTEMS IN DATA CENTERS	2
1 PROBLEM STATEMENT	2
2 TECHNICAL OBJECTIVES	2
3 TECHNICAL INFORMATION ON THE CHARACTERISTICS OF MODULAR, SCALABLE COOLING SYSTEMS AND SERVERS.....	3
3.1 MODULAR, SCALABLE, OVERHEAD COOLING SYSTEM	3
3.2 SERVERS	3
3.3 SERVER POWER MANAGEMENT.....	4
3.4 EQUIPMENT LOCATION	4
4 MEASURED PARAMETERS.....	6
5 DATA ACQUISITION SYSTEM.....	7
6 TEST PROCEDURES AND OPERATING CONDITIONS.....	7
7 PERFORMANCE METRICS FOR MODULAR COOLING	8
8 SUMMARY OF FINDINGS AND CONCLUSIONS.....	10
9 ACKNOWLEDGEMENTS.....	14
10 REFERENCES.....	14

Evaluation for Modular, Scalable Overhead Cooling Systems 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.

Naturally, the trend toward higher power density resulting from current and future generations of servers has, in the meanwhile, 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 1.

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 overhead cooling system. The system was tested in a hot/cold aisle environment without separation, or containment or the hot or cold aisles. The scope of this report 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.

The information generated from this testing when combined with a concurrent research study to document the 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 modular, scalable cooling systems and servers

3.1 Modular, scalable, overhead cooling system

Cooling system. The cooling system in this study is a pumped refrigerant cooling system designed for cooling data center IT equipment racks. This system was configured for a total cooling capacity of 80 kW, and consisted of one pump unit and eight cooling modules. A typical system consists of one pump unit and up to 16 cooling modules, and can provide up to 160 kW of cooling. Multiple cooling systems may be used for higher cooling demands.

Pump unit. The pump unit includes a pump that circulates the refrigerant and a refrigerant-to-water heat exchanger. This heat exchanger transfers heat from the pumped refrigerant into the chilled water that circulates through the pump unit and back to the central chilled water source.

Cooling module. Each cooling module contains two fans and an air-to-refrigerant heat exchanger. Hot air from the hot aisle was pulled into the cooling module and heat was transferred into the pumped refrigerant through the air-to-refrigerant heat exchanger. Cool air is then provided to the cold aisle in front of the computer racks.

The water flow rates in the pump unit as well as the airflow rates in the cooling modules were controlled with constant-speed pump and fans. Regardless of the cooling load, the power consumption of the cooling system was largely fixed. For the servers with varying loads in this study, some experiments were included with only one single fan running under partial load.

3.2 Servers

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". Due to constraints in data center space for the study, and the high capital cost for servers, 320 servers were selected to be stacked in eight server racks. As a result, the maximum nominal load per server rack was designed to be 10 kW. 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 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 various heat load levels in the test plan, the number of servers running at full load per rack were as follows: 10 kW/rack - 33 servers; 7.5 kW/rack - 25 servers; 5 kW/rack - 17 servers; 2.5 kW/rack - 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 server racks and cooling modules in this study were located in the northwest corner of the data center. The space containing all server racks and cooling modules has floor dimensions of 8’6” by 22’, and was separated from the rest of data center by an array of transparent curtains surrounding the server racks and eight cooling modules. Eight 42U server cabinets, each with the size of 29.5”by 42” by 78,” were oriented in a single row, with a common cold aisle (42” wide) on the inlet side of the servers, and a common hot aisle (18” wide) on the outlet side of the servers. Air from the surrounding CRAC units in the vicinity was intended to be blocked off using the curtains. The asymmetric arrangement between cold- and hot-aisles allowed sufficient cold-aisle space for normal air movement.

One overhead cooling module (sizing 14 inches high, 24 inches wide, and 39 inches long) was positioned at 4.5 inches above the top of each of the eight server cabinets. The overhead cooling modules are designed to draw air from the hot aisle, passing it through the air-refrigerant heat exchanger and then directing cold air to the cold aisle of the servers. The cold air delivered overhead to the cold aisle is then moved from the inlet of

the servers through the rack to the hot aisle where the process repeats itself. Figure 1 and Figure 2 illustrate the general equipment layout in this study.

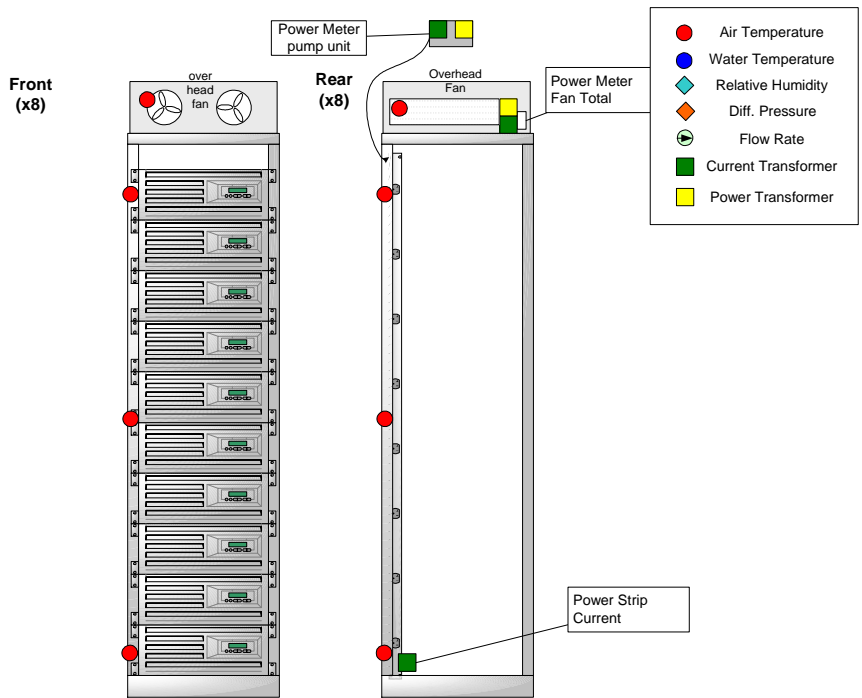


Figure 1. Equipment layout – Front view




Figure 2. Equipment layout – Top view

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.
 - 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
- Pump unit entering and leaving chilled water temperatures

Chilled water temperatures in the supply and return pipes were measured by installing two surface mounted temperature sensors on the supply and return chilled water lines. Insulation was wrapped around the pipes to eliminate the influence of the surrounding environment to the pipe sensors.
- Pump unit chilled water flow rate
 - The flow rate was calculated to be within $\pm 5\%$ accuracy using the pressure drop monitored and a vendor supplied conversion equation. Pressure drop was monitored across a balancing valve on the supply side of the chilled water connected to the refrigerant-water heat exchanger.

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 IT equipment power; total central chilled water plant power; primary chilled water flow and chilled water temperatures (supply and return); chilled water supply/return differential pressure.

5 Data acquisition system

Data collection for the test environment was performed using a commercially available data acquisition system [3]. 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 pumping system, the fans, and one 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 temperature 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 an ultrasonic flow meter, pressure transducers, and RTDs. The cooling hardware itself did not have communication capabilities.

6 Test procedures and operating conditions

The supply water temperature to the pump unit used in this evaluation was 45°F, which was the actual chilled water temperature from the central cooling system.

Selected operating conditions were designed by combining various server loads (25% to 100%) and various inlet air temperatures in this study, and are presented in the following table.

Table 1 Set Points for Test Conditions

Inlet air temperature set point (F)	Targeted total server load (kW) for 8 servers racks	Various server loads (kW) per rack			
		100%	75%	50%	25%
68	80	10.0	7.5	5.0	2.5
72	80	10.0	7.5	5.0	2.5
76	80	10.0	7.5	5.0	2.5
80	80	10.0	7.5	5.0	2.5

- For each desired rack power load, temperature control measures were taken to adjust inlet air temperatures at discrete set points, i.e., 68°F, 72°F, 76°F, and 80°F, respectively.

- It was noted that the Modular Cooling System 1 - overhead cooling units used in this study are designed to maintain a Minimum Room Temperature and are not specifically designed to provide direct set point control over inlet air temperature. Therefore, the inlet air temperature being delivered to the cold aisle is designed to float based upon the rack power or heat load of the rack. For this reason, it was difficult to achieve the exact set points for inlet air temperatures as contemplated in the test plan, especially during part load since the heat rise across the servers in the rack was not high enough to reach the low end of the lower inlet air temperature set points since the units are not designed to heat the air. It was also noted that although initial trial tests followed the rule-of-thumb provided by the vendor that it took about 30 minutes to achieve steady-state air temperature, we typically allowed from 45 minutes to up to two hour duration to enable the system to reach steady state for each operating condition based upon the field-testing experience gained in this study.
- Change the server load from full (i.e., 100%) to 75%, 50%, and 25% of the defined nominal capacity per rack (i.e., approximately 10 kW/rack). Repeat and record the measurements for each operating condition.

7 Performance metrics for modular cooling

In order to characterize thermal performance of the cooling system, we used the ratio of cooling (in kW) transported by the cooling module to the total power demand for the operation of the module (refrigerant pump and fan), defined as “coefficient of performance (COP).” The COP is unitless, with a higher value representing higher efficiency for the cooling module. The COP can be calculated under achieved operating conditions (a range determined by inlet air temperature and server load).

In this evaluation, the total power demand is the pump and fan power required to transport the required refrigerant and airflow from the cooling system, while the heat removed is equivalent to the cooling provided by the cooling system.

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

Where

Cooling is the amount of cooling power removed from the data center by the cooling system, and

P_{total} is defined as the total power demand for the cooling system to do so.

Because there were fans and refrigerant pumps in the Cooling System 1, the total power demand was for pumping and air-circulating requirement in this study.

Total power demand for the module can be measured as follows:

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

Where

$P_{fan\ power}$ is the total measured fan power demand of the cooling modules

$P_{pump\ power}$ is the measured pump power demand of the pump unit

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

Where

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

$Cooling$ is the heat transfer (in kW) from the chilled water plant via the heat exchanger in the pumping module.

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

Q : Averaged chilled water flow rate measured in gallon per minute

The pressure difference between supply and return across the valve was used to estimate the actual water flow rate across the valve. The pressure-to-flowrate conversion formula was based upon the calibration provided by the balancing valve vendor, with an error bar of $\pm 5\%$.

$$Q = 40.2\sqrt{\Delta P}$$

where ΔP is the pressure difference of the chilled water flow in psi.

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

ΔT_w : Measured water temperatures rise, in °F

Therefore,

$$Cooling = 0.1467Q \Delta T_w$$

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

Because we were focused on modular cooling system in this evaluation, 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 overhead 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.

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

$$MSE = \frac{3.52}{COP}$$

where

Module System Efficiency (*MSE*): ratio of total cooling power to the cooling transported by the module, 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 module's power index, *PI*, is the ratio of power demand for the cooling system to computer load under selected operating conditions. A higher value of the power index indicates higher cooling energy demand for the cooling system at a given server load.

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

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 and 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 10kW and depending on the server load set points. The control 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.

Based upon the three temperature measurements at the server rack inlet and outlet, average top, middle and bottom temperatures were calculated for rack. These temperature

measurements were used to calculate the average, maximum, minimum, and standard deviation among all eight racks in operation during this study, as shown in Table 2.

Table 2 Statistics for inlet air temperatures and average outlet air temperatures

Average Inlet Air Temperature (F)	Maximum Inlet Air Temperature (F)	Minimum Inlet Air Temperature (F)	Standard Deviation of Inlet Air Temperature (F)	Average Outlet Air Temperature (F)	Range (difference of max and min) (F)
74.6	75.9	72.4	1.0	102.4	3.5
73.1	75.8	68.5	3.0	100.6	7.4
74.1	78.1	64.3	5.2	100.4	13.8
76.4	81.2	64.9	6.9	103.4	16.3
73.3	76.4	67.1	3.6	99.9	9.3
71.9	73.3	69.1	1.4	94.8	4.2
72.1	75.2	65.1	4.1	96.7	10.0
79.7	80.9	77.7	1.2	105.1	3.1
71.5	73.2	69.9	1.2	95.0	3.2
65.6	66.4	65.1	0.5	86.4	1.4
65.3	66.5	64.8	0.6	86.6	1.7
65.3	66.5	64.9	0.6	86.6	1.6
74.7	76.2	73.7	1.0	98.7	2.5
68.5	69.4	68.0	0.5	82.4	1.4
66.0	66.8	65.5	0.6	84.0	1.4
68.3	69.0	67.8	0.4	85.1	1.3
71.1	72.0	69.7	0.8	86.9	2.3

The distribution of inlet air temperatures in this study appeared to be non-uniform, from the bottom to the top of the server racks, and from rack to rack. For example, the difference between maximum and minimum inlet temperatures ranged from less than 2°F up to 14°F, with the standard deviation ranging from less than 0.4°F up to 7°F. The non-uniformity among the eight racks tested was particularly evident when inlet air temperatures were intended to be controlled within the higher end of the set points shown in Table 1.

The actual operating conditions achieved in this evaluation are shown in Table 2. The overall coverage of operating conditions ranged from 65°F to 80°F for inlet air temperatures and the server loads ranged from 2.7 kW/rack to 9.8 kW/rack.

Table 3 Actual Test Conditions and Results

Server Load per Rack (kW/rack)	Average Inlet Air Temperature (F)	Average Outlet Air Temperature (F)	Chilled Water Flowrate (gpm)	Chilled water temperature difference (F)	Cooling from Chilled Water (kW)	Fan Power (kW)	Pump Power (kW)	Total Power for Cooling Module (kW)	Total Server Power (kW)	COP	MSE (kW/ton)	Power Index
9.6	75	102	46.6	8.9	61.5	1.5	0.9	2.4	76.8	26.2	0.13	0.03
9.6	73	101	61.6	8.5	78.6	1.5	0.9	2.4	77.0	32.9	0.11	0.03
9.7	74	100	36.5	13.3	70.9	1.5	0.9	2.3	77.5	30.3	0.12	0.03
9.8	76	103	23.9	18.3	66.1	1.5	0.9	2.3	78.6	28.2	0.12	0.03
9.7	73	100	60.0	9.2	82.9	1.5	0.9	2.3	77.3	35.5	0.10	0.03
7.5	72	95	30.6	13.5	60.4	1.5	0.9	2.4	60.1	25.1	0.14	0.04
7.6	72	97	24.3	14.6	53.5	1.5	0.9	2.4	60.7	22.4	0.16	0.04
7.6	80	105	10.4	24.7	59.6	1.5	0.9	2.4	61.0	25.2	0.14	0.04
4.9	71	95	26.5	9.9	38.8	0.7	0.9	1.6	39.6	24.0	0.15	0.04
4.8	66	86	59.8	6.9	62.0	1.5	0.9	2.3	38.5	26.4	0.13	0.06
4.8	65	87	42.6	7.6	50.4	1.5	0.9	2.3	38.4	21.5	0.16	0.06
4.8	65	87	44.9	7.1	55.9	1.5	0.9	2.4	38.3	23.8	0.15	0.06
5.0	75	99	52.7	18.0	43.0	0.7	0.9	1.6	40.0	26.7	0.13	0.04
4.8	68	89	9.9	13.3	45.8	1.5	0.9	2.3	38.7	19.6	0.18	0.06
2.7	66	84	22.4	-	-	0.8	0.9	1.7	21.8	-	-	0.08
2.7	68	85	0.0	-	-	0.8	0.9	1.7	21.7	-	-	0.08
4.8	71	87	0.0	-	-	0.8	0.9	1.7	38.5	-	-	0.04

Table 3 shows the results from the tests performed at the facility including: server load; average inlet/outlet air temperatures monitored at three different heights for all eight server racks; cooling delivered by the modules; power demand of the cooling systems, and three performance metrics in this study: COP; module system efficiency; and PI.

It is clear that different IT equipment operation 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 COP, module cooling efficiency (kW/ton), and power index for modular Cooling System 1.

Generally, modular pump power demand remained stable (mostly around 0.9 kW), while fan power demand exhibited a slightly higher demand when both overhead unit fans were on (1.5 kW) compared to the fan power when one fan was on. Overall, the total power demand for the modular cooling was very constant for a given operation condition, i.e., 2.4 kW with dual fans on and 1.6-1.7 kW with one fan for each of the eight cooling modules.

Under similar server loads, the COP of cooling system 1 tended to increase with the decrease in inlet supply air temperature; while the cooling system’s MSE (kW/ton value) tended to increase with the increase in inlet supply air temperature. The magnitude of MSE ranged between 0.1 and 0.2 kW/ton, lower than a normal computer room air handling (CRAH) unit. The PI value appeared to change little under similar server loads with various inlet air temperatures in this study.

As expected for a system that consumes electric energy at a relatively constant rate, the COP values of the modular cooling system increased with the increase in server load;

indicating higher cooling efficiency with higher rack power loads. MSE values show a reversed trend, also indicating higher cooling efficiency at higher server power loads. Similarly, the PI values increased with the decrease in rack power load, indicating that higher cooling system power per unit of server heat demand was required under lower server loads.

The test results show that this cooling system was generally capable of providing cooling needed to achieve various inlet air temperatures under various server loads pre-designed for the study; however, there were challenges in the course of controlling and maintaining inlet air temperature within a desired range for some of the set points used in this evaluation. Specifically, the challenge in controlling cooling conditions was most evident at the high end of the selected set point temperatures in this study.

Overall, the calculated COP values ranged from 20 up to 36; MSE (kW/ton) values ranged from 0.11 up to 0.18kW/ton; and PI ranged from 0.03 to 0.08. Modular Cooling System 1 exhibited an energy efficiency level better than traditional CRAH units under the selected operating and environmental conditions (e.g., pre-defined rack power and inlet air temperature) that were achieved in this study. The findings from this study indicate that by implementing overhead modular 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 considered for successful implementation of modular 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. In addition, the overall energy demand will be also be affected, to some extent, by the individual CRAH units or other modular cooling units installed within the data center. In this evaluation, the cooling system was operating at the chilled water temperature of approximately 45°F provided by the central chilled water plant.

It would be useful to understand the cooling effectiveness by coupling modular cooling units with the chilled water plant supplying cooled water of various temperatures. This study recommends that the cooling performance and energy efficiency of Cooling System 1 be evaluated when operating with higher supply water temperature up to the vendor’s recommended maximum level.

In addition, due to testing constraints, evaluation of cooling performance was not performed for Cooling System 1 at its maximal design load cooling capacity (i.e., 160 kW for 16 modular overhead cooling units). Since the heat exchanger/pump system has constant power demand and is designed for a maximum of 160kW, future tests may be conducted to test more typical design operating cooling load, e.g., 80% or 128kW, while using the vendor’s cold aisle containment system to better understand how the system performs both operationally and from an energy efficiency perspective.

In order to further improve energy efficiency of the modular overhead cooling system via reducing energy use by fan and pump operation, it is recommend that the vendor consider controllability to provide variable flow rates on the heat exchanger/pump system and to

automatically operate the overhead units in a single or dual fan mode in future generations of the product.

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. Based on the magnitudes of the performance metrics developed and evaluated in this study, it is clear that this modular cooling system can be more efficient than the typical CRAH units widely used in traditional data centers. This being said, however, it would be premature to directly compare this modular overhead cooling system with those of other similar modular, scalable cooling systems because of their differences in actual operation conditions and optimal design loads that could affect actual efficiency outcomes from the tested operation. In addition, it is recommended that the reader of these individual reports 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 the data center.

In closing, 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, the local weather conditions where the datacenter is located, 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, Inc.
<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.