LBNL-XXXXX

ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

Direct Liquid Cooling for Electronic Equipment

Henry Coles and Steve Greenberg
Environmental Energy Technologies Division

March 2014

DISCLAIMER

This report was prepared as the result of work sponsored by the California Energy Commission. It does not necessarily represent the views of the Energy Commission, its employees or the State of California. The Energy Commission, the State of California, its employees, contractors and subcontractors make no warrant, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the uses of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the California Energy Commission nor has the California Energy Commission passed upon the accuracy or adequacy of the information in this report.

ACKNOWLEDGEMENTS

This work was supported by the California Energy Commission's Public Interest Energy Research Program and by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Thanks to key partners Asetek (San Jose, California), Cisco Systems (San Jose, California) and Intel for the cooling systems, computer equipment, and processor components used for the demonstration. Specific appreciation goes to Steve Branton, Anderson Tsai, and Steve Empedocles (Asetek); Roy Zeighami (Cisco Systems); and Winston Saunders (Intel).

Another key partner was Joulex (acquired by Cisco during the project), which provided the software and programming support used to collect and store the data from the varied sources. Specific appreciation to Scott Paisley and Oliver Stöneberg. Thanks to Server Technology, Inc. who provided power measurement and distribution equipment.

Syska Hennessy Group (San Francisco, California) provided data center engineering consulting and modeling to support the energy-efficiency analysis. Thanks to Vali Sorell, V.P., and Ian Kucma, the lead Syska Hennessy Group engineer.

This project would not have been possible without the tireless support provided by colleagues working in the High Performance Computing Services group at Lawrence Berkeley National Laboratory, headed by Gary Jung. A special thank you to Susan James for her support of IT equipment and network configuration and issue resolution during the project. Thank you to Yong Qin who supported the project with software development required for data queries, data conversion, and analysis. Additional thanks to Karen Fernsler and Bernard Li, and many thanks to Mike Magstadt, who supported the effort by allowing us to use a Lawrence Berkeley National Laboratory data center for this project.

ABSTRACT

This report documents a demonstration of an electronic-equipment cooling system in the engineering prototype development stage that can be applied in data centers. The technology provides cooling by bringing a water-based cooling fluid into direct contact with high-heat-generating electronic components.

This direct cooling system improves overall data center energy efficiency in three ways:

- High-heat-generating electronic components are more efficiently cooled directly using water, capturing a large portion of the total electronic equipment heat generated. This captured heat reduces the load on the less-efficient air-based data center room cooling systems. The combination contributes to the overall savings.
- The power consumption of the electronic equipment internal fans is significantly reduced when equipped with this cooling system.
- The temperature of the cooling water supplied to the direct cooling system can be much higher than that commonly provided by facility chilled water loops, and therefore can be produced with lower cooling infrastructure energy consumption and possibly compressor-free cooling.

Providing opportunities for heat reuse is an additional benefit of this technology. The cooling system can be controlled to produce high return water temperatures while providing adequate component cooling.

The demonstration was conducted in a data center located at Lawrence Berkeley National Laboratory in Berkeley, California. Thirty-eight servers equipped with the liquid cooling system and instrumented for energy measurements were placed in a single rack. Two unmodified servers of the same configuration, located in an adjacent rack, were used to provide a baseline.

The demonstration characterized the fraction of heat removed by the direct cooling technology, quantified the energy savings for a number of cooling infrastructure scenarios, and provided information that could be used to investigate heat reuse opportunities.

Thermal measurement data were used with data center energy use modeling software to estimate overall site energy use. These estimates show that an overall data center energy savings of approximately 20 percent can be expected if a center is retrofitted as specified in the models used.

Increasing the portion of heat captured by this technology is an area suggested for further development.

Keywords: direct liquid cooling, direct cooling, direct-chip cooling, liquid-cooled computer, liquid-cooled chip, data-center liquid cooling, data center efficiency

TABLE OF CONTENTS

Acknowledgements i

ABSTRACT ii

TABLE OF CONTENTS iii

LIST OF FIGURES v

LIST OF TABLES v

EXECUTIVE SUMMARY 1

Introduction 1

Project Purpose 1

Project Results 2

Project Benefits 3

CHAPTER 1: Introduction 4

1.1 Summary 4

1.2 Demonstration Goals 6

1.3 Participants 6

CHAPTER 2: Methods 7

2.1 Introduction 7

2.2 Goals 7

2.3 Demonstration Process 7

2.4 Required Data 9

2.5 Test Setup 10

2.5.1 Physical Layout 10

2.5.2 Equipment Descriptions 12

2.5.3 Data Collection 13

2.5.4 Test Procedure 15

2.5.5 Calculation Process 16

2.6 Data Center Energy Use Models 18

CHAPTER 3: Results **22**

3.1 Maximum Heat Removed Using Direct Cooling 22

3.2 Heat Reuse from Direct Cooling 24

3.3 Site Energy Consumption 25

CHAPTER 4: Conclusions and Recommendations **29**

GLOSSARY **30**

REFERENCES **31**

APPENDIX A: Test Plan Parameter Combinations **1**

APPENDIX B: Lytron CDU Plumbing Detail **1**

APPENDIX C: Electrical Power and Btu Meters **1**

APPENDIX D: Air-Server Power Correction **1**

APPENDIX E: Statistical Analysis of Fraction Cooled **1**

APPENDIX F: Thermal Test Results **1**

APPENDIX G: Data at 100%, 50%, and Idle Server Power **1**

APPENDIX H: Modeling Results: 100%, 50% and Idle Server Power **1**

LIST OF FIGURES

Figure 1-1: Schematic: Direct Cooling System 5
Figure 2-1: View Of Rack Containing 38 Cisco Servers Equipped With Asetek Direct Cooling Technology - Lytron CDU Shown In Foreground 11
Figure 2-2: Demonstration Schematic Layout 12
Figure 2-3: Schematic: Data Collection 15
Figure 2-4: Base Case Configuration Cooling Infrastructure 19
Figure 2-5: Cooling Tower and Cooling Tower With Chiller Boost Cooling Infrastructure Case Model 20
Figure 2-6: Dry Cooler And Dry Cooler With Chiller Boost Cooling Infrastructure Case Model 20
Figure 2-7: Direct Cooling Provided by Existing Chiller Infrastructure Case Model 21
Figure 3-1: Maximum Heat Capture Vs. Water Supply Temperature For Three Sever Power Levels 23
Figure 3-2: Heat Captured Vs. Return And Supply Water Temperature 24
Figure 3-3: 100% Server Power Case Annual Energy Use Estimates 26
Figure 3-4: 50% Server Power Case Annual Energy Use Estimates 27
Figure 3-5: Idle Server Power Annual Energy Use Estimates 28
Figure B-1: Lytron CDU Piping Details B-1
Figure C-1: Installed Electrical and Btu Metering C-1
Figure D-1: Composite Server Power D-1
Figure D-2: Air-Server Power Front Panel Air Inlet Temperature Correction Graph, Using Data from Figure D-1 D-1
Figure G-1: Time Series Line Graphs - 100% Server Power Test Data G-1
Figure G-2: Time Series Line Graphs - 50% Server Power Test Data G-2
Figure G-3: Time Series Line Graphs - Idle Server Power Test Data G-3

LIST OF TABLES

Table 2-1: Collected Data Descriptions 14
Table 2-2: Parameter Combinations Used for Site Energy Use Estimates 21
Table A-1: Test Plan Parameter Combinations for Servers at Idle Power Level A-1
Table A-2: Test Plan Parameter Combinations for Servers at 50% Power Level A-1
Table A-3: Test Plan Parameter Combinations for Servers at 100% Power Level A-2
Table F-1: Thermal Performance Test Results F-1
Table H-1: Model Results For 100% Server Power H-1
Table H-2: Model Results For 50% Server Power H-2
Table H-3: Model Results For Idle Server Power H-3

EXECUTIVE SUMMARY

Introduction

Data centers in the United States currently consume approximately 2 percent of the nation's electrical energy. In data centers, a large part of that electrical energy, often 50 percent or more, is consumed by the power distribution and cooling systems. The energy consumed by the cooling systems is often on the order of two or more times that of the losses in the power distribution systems. Therefore energy savings from improving cooling system efficiency are significant.

The technology demonstrated in the project is a system that provides cooling by bringing cooling water directly to heat-sensitive electronic components. The concept of this technology is not new; it has been and is currently being applied to high-performance computers and super computers. However, providing this cooling method at a price compatible with the commercial data center market is relatively new.

Overall data center energy savings are achieved by applying this technology to cool a significant fraction of the heat produced by the total electrical energy supplied to the electronic equipment. The technology captures heat at the electronic component and transfers it directly to a liquid, avoiding a less-efficient heat transfer method using air. The heat captured using the technology reduces the cooling load on the less-efficient data center room air conditioning systems, providing a net reduction in energy use for the data center. A higher fraction of heat captured provides lower overall data center energy use. In addition, this technology can provide the required cooling using high-temperature water—much higher than typically found in data centers. Use of the high-temperature supply water results in an even higher return water temperature, opening opportunities for heat reuse.

This demonstration explored these capabilities: fraction of heat captured, basic capabilities as applied to heat reuse, and overall data center energy use.

Project Purpose

New technologies or technologies applied with a different method or in a new market are frequently encountered. This is especially true of cooling as applied to the data center market. When a technology is presented by the manufacturer along with energy-efficiency claims, data center owners are interested in confirming the validity of those claims. The purpose of this demonstration was to quantify the data center energy-efficiency capabilities of this new product prototype using measurements taken in an active data center.

In addition, the project documented methods for gathering empirical data and the application of those data for estimating overall data center energy, to assist others with future studies of the same or similar technologies.

Project Results

The results are divided into two sections:

- Heat Capture
  - For maximum heat removed
  - For heat reuse applications
- Data Center Energy Use Savings

Heat Capture: Maximum Efficiency

The results for maximum heat capture were widely variable. The heat captured is expressed as a percentage of the power supplied to the electronic equipment by their power cords.

The best results ranged from 67 percent to 75 percent. These maximum heat capture results corresponded with low supply water temperatures (15°C–20°C; 59°F–68°F) and were not significantly affected by the electronic equipment power level. The maximum heat captured when higher supply water temperatures were used were much different, and varied considerably as a function of equipment power level. For example, the heat captured at the server idle power level was 35 percent using a 30°C (86°F) supply water temperature. The heat captured for the 50 percent and 100 percent power level was 57 percent and 59 percent, respectively, using 30°C supply water temperature.

Heat Capture: Heat Reuse

The data for heat captured considering heat reuse applications consisted of measurements using a range of supply water temperatures and reducing the water flow rate to achieve a return water temperature that could be supplied to a heat-reuse application.

An interesting phenomenon was observed: at a constant IT equipment power level, the heat captured was little affected by the supply water temperature, but the heat captured was a strong function of the return water temperature. The percentage of heat captured was close to a linear function of the return water temperature, with higher temperatures resulting in lower heat capture percentages.

Overall Energy Savings

The overall data center energy use saving estimates were made using models constructed with Romonet data center energy use simulation software. A number of simulation models were constructed using the software, including the addition of a dry cooler or cooling tower to support the direct cooling system. The case of using the existing chilled water supply was also included in the models. A base case that used the measured data from the unmodified servers was used to estimate the overall data center energy use savings.

The results were grouped by the three IT equipment power level categories: 100 percent, 50 percent, and idle.

The energy saving estimates using the 100 percent power-level measurements comparing the base case and direct cooling cases were similar, ranging from 17 percent to 23 percent.

The 50 percent power level created by the artificial load software was considered to be representative of the load consistent with a high-performance computer data center. The results for the 50 percent power level were similar to those for the 100 percent power case but were 2 to 5 percentage points lower, ranging from 14 percent to 20 percent overall savings. An interesting infrastructure case, using the existing chilled water, resulted in an overall savings estimate of 16 percent. Using existing chilled water may not provide the best operational savings. However, the potential for low-cost implementation in a retrofit (utilizing existing systems) could offset the higher operational costs, thereby providing a lower total cost of ownership.

The energy savings estimates for the idle power cases were significantly lower, 9 percent compared to the savings for 50 percent and 100 percent power cases mentioned above.

The results did not include the data center ambient temperature as a variable. The measurements and results were normalized as needed to a server air inlet temperature of 28°C (82°F). The heat captured as a function of server air inlet temperature is an area for further study.

The IT equipment equipped with this technology consumed less energy compared to the unmodified version in all cases. The savings in IT energy was a significant part of the overall energy savings. When considering implementation of this technology in a new or existing data center, it is suggested that estimates of IT energy savings be supported by measurements using actual hardware and software.

Quantifying the reliability of this cooling system was not a goal of this demonstration; however, it should be noted that no water leaks were observed during the more than six-month demonstration.

Project Benefits

This report documents an empirical method for determining the fraction of server heat that is removed using this cooling technology, thereby reducing the heat load on existing and less-efficient computer room air conditioning systems.

A number of cooling infrastructure retrofit modification scenarios were investigated, using modeling to explore the change in overall data center energy use if this technology were to be implemented in a commercial data center. Because the climate is a model input, it is important to note that San Jose, California, climate data were used in the models. The modeled cooling infrastructure modifications were selected as reasonable for a retrofit scenario for data centers located in California.

Heat-capture performance results and overall energy use estimates included in this report will help data center owners make informed decisions regarding a change to this type of hybrid-cooling equipment for electronics in their data center.

CHAPTER 1: Introduction

1.1 Summary

Data centers in the United States are estimated to be consuming approximately two percent of the electrical energy consumed countrywide (Koomey 2011). Approximately one-half of this energy is used to provide cooling and other support (also know as *infrastructure energy*) for the electronic equipment. The cooling portion of this support historically and currently is the dominant contribution to the total infrastructure energy requirements.

In recent years, new products, improved controls, and expanded operational limits have been demonstrated to provide energy savings. Examples include improving airflow management, improving computer room air handler (CRAH) or computer room air conditioning (CRAC) automation (Coles 2012), or providing close-coupled cooling at the rack level.

These design and operational improvements are effective, but they do not take full advantage of the heat transfer opportunity available at the hot surfaces on electronic components, such as those for processing and memory. Temperatures found at these surfaces can reach 70°C (158°F) or higher while functioning within specified limits.

Commercial data centers contain and continue to purchase IT equipment that is cooled by air, using heat sinks and internal fans to cool the fins on the heat sinks and other components. This method does not take advantage of other, more direct, heat transfer approaches using liquids.

Cooling electronic components that use a liquid applied directly to or near the surface of the chip is not a new technology. This approach has been used in past, and now almost exclusively on mainframe or high-performance computers commonly found in supercomputer facilities.

In recent years, economical versions of direct-cooling using water have been developed and sold to the personal computer market for those customers wishing to maximize performance. Now a version of the technology, with modifications, is available for the commercial server market. A version of this new application, developed by Asetek, was used for the demonstration presented in this paper.

The term *direct cooling* will be used to describe the heat transfer from electronic components provided by the Asetek system starting at the chip and ending in the building chilled water system via the water-to-water heat exchanger referred to as the *RackCDU™*.

Direct cooling provides a more-efficient method to transfer the heat from these hot components to the building chilled water loop and then outside with very little additional energy, compared to transferring the heat first to air and later to the building chilled water system.

In addition, in a direct cooling system, the water temperature returning after cooling the IT equipment is much higher than typically found in data centers, and provides more opportunity

for heat reuse or the ability to reject this heat to the atmosphere with a dry cooler, thereby eliminating the requirement of a cooling tower or chiller plant in most climates.

A dry cooler system is nothing more than a water-to-air heat exchanger, with fans to provide airflow. Pump energy also is needed to move the water between the data center and the dry cooler. The combination of the dry cooler fans and pumping consume a small fraction of the energy needed to operate a typical chilled water plant.

If climate or physical space prohibits using a dry cooler, in some cases the existing cooling infrastructure can be modified to support the direct cooling processes and still provide significant overall savings.

Figure 1-1 shows the location of the system components and how they work with each other.

Integrated pump and cold-plate assemblies absorb heat from the central processing unit (CPU). The memory dual in-line memory modules (DIMMs) are cooled by transferring heat to a manifold (in contact using heat transfer tape) carrying the cooling water. The cooling water supply and return paths are provided by a set of flexible tubes for each server. The heat collected is transferred using the tube set to the facility cooling support by means of a cooling distribution unit (CDU), as shown in Figure 1-1.

Some of the cooling for the server is not provided by the direct cooling system; there are a number of electronic components inside the server that still need to be air-cooled, using fans inside the server. However, the airflow requirement is reduced, and therefore the number of fans and their speed is reduced, providing a significant reduction in the energy consumed by the server for the same processing load. This server fan energy reduction accounts for most of the reduction in server energy use.

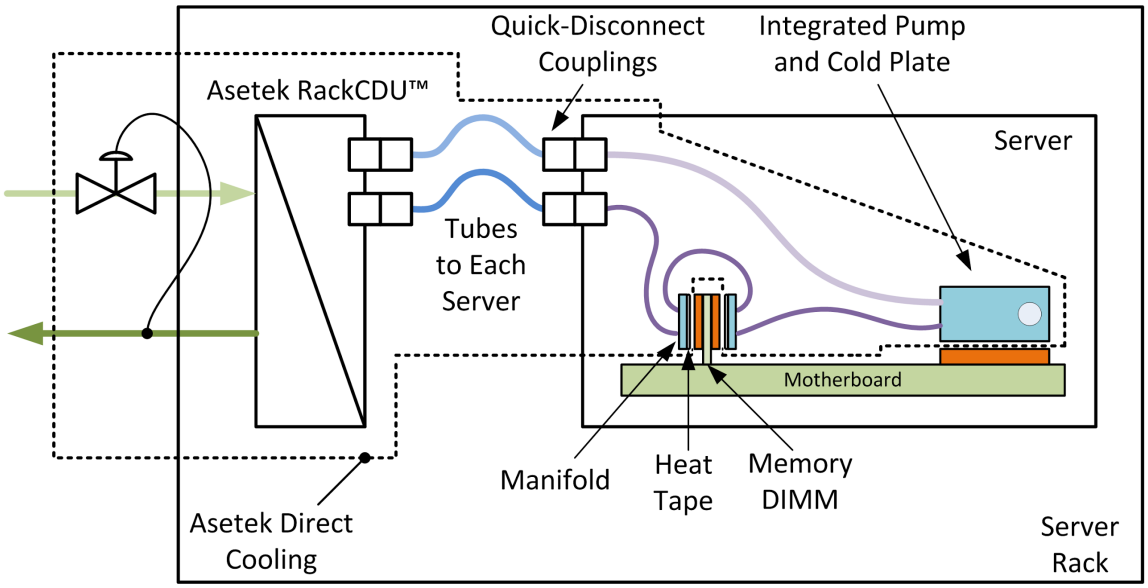

Figure 1-1: Schematic: Direct Cooling System

1.2 Demonstration Goals

The goal of this demonstration was to determine the thermal performance of this prototype cooling system and its potential impact on overall data center energy use.

The project hypothesis: A data center that substitutes the prototype cooling system for a conventional air-cooled method will consume less energy.

The analysis used a two-step approach: (1) measure the fraction of power supplied to the servers that is captured as heat by the prototype cooling system, and (2) use these fractions to calculate estimates using models comparing a server equipped with a stock cooling configuration (air cooled) to a prototype cooling system (direct-water cooled).

1.3 Participants

A number of organizations were involved in this demonstration in addition to those involved at Lawrence Berkeley National Laboratory (LBNL). Asetek Inc. provided the complete RackCDU cooling system, including retrofit of 38 servers and technical support. Cisco provided 40 servers and technical support. Intel provided the CPUs for the 40 Cisco servers. Joulex Inc. (which became part of Cisco during the project) provided data collection software and programming. The Syska Hennessy Group provided data center modeling using Romonet software and also provided technical guidance. Sever Technology, Inc. supplied two-rack power distribution units, enabling individual server power metering.

CHAPTER 2: Methods

2.1 Introduction

A series of energy-use related tests were completed using a set of modified and unmodified IT equipment. The IT hardware equipped with the prototype cooling technology consisted of 38 modified Cisco C220 M3 servers loaded in a standard data center server rack. Two standard unmodified servers located in the bottom of an adjacent rack were not equipped with the Asetek cooling technology, and were used to establish a base case.

As mentioned above, the testing and analysis consisted of two steps: (1) measure the thermal performance (amount of heat captured), and (2) estimate data center energy savings using the data from the first step.

The data for calculations to complete the first step were obtained by measurement.

The results for the second step were obtained by entering the results (the fraction of heat captured) from the first step into a number of data center models created to simulate overall data center energy use for various retrofit cooling infrastructure configurations.

2.2 Goals

The demonstration goal was to explore the energy-savings potential of the use of this technology in data centers. The method was documented to assist others with future evaluations of equipment that use the same or similar IT equipment cooling designs.

This demonstration was not an attempt to quantify results using exacting statistical methods. The results are presented in graphical form to help the reader see trends that could be confirmed later with more-specific testing.

2.3 Demonstration Process

This technology provides a replacement cooling method, for some components, internal to the IT equipment. The standard cooling method uses exclusively the flow of air provided by internal fans to keep electrical components within acceptable operating temperatures. The technology used in the project replaces the air cooling provided for high-heat generating components with cooling provided by a liquid. Therefore, IT equipment using this technology has two heat removal paths (air and liquid) compared to the standard IT equipment having only one. This replacement cooling technology is referred to as the *direct cooling system*.

Estimates of the energy savings potential were performed by analyzing two data center scenarios: (1) a center filled with standard servers, and (2) a center filled with servers retrofitted with the direct cooling system.

Energy savings cannot be measured, but must be a comparison of two estimates: the energy use for the base case (standard design) and the energy use for the retrofit case. The energy-use estimates were completed for the base case and the direct-cooling case in two steps:

Step 1: Evaluate Stand-Alone Thermal Performance

Determine the fraction of the heat captured by the direct cooling system for parameter combinations covering a broad range of conditions. This provided the necessary information used in Step 2.

Step: 2: Estimate Data Center Site Energy Consumption

The site energy use was estimated using the results from Step 1. Modeling results using the data center configurations are described below.

Simple comparisons of the base case and retrofit case results from Step 2 are made to obtain an idea of the energy savings that might be expected.

Parameters Varied

The parameters varied during the tests were: supply water temperature, supply water flow rate (which affects the return water temperature), and server power. The results from varying the supply water temperature provides data for modeling different cooling infrastructures. The performance associated with higher return water temperatures was used to investigate the possibilities of heat reuse. The server power was varied to investigate the effects of different levels of server utilization.

In this demonstration, the liquid going from the RackCDU to the servers was provided by Asetek. This fluid is mostly water, but it also had anti-corrosion and freeze protection additives. Hereinafter, the liquid will be referred to simply as *water.*

The parameters *supply water temperature* and *return water temperature* refer to the water flowing between a Lytron CDU (foreground in Figure 2-1) and RackCDU indicated in Figure 2-2 as "cooling water supply/return." The Lytron CDU contains a water-to-water heat exchanger, pump and controls. This CDU was used to adjust the supply water temperature.

The lowest supply water temperature tested for this demonstration was 15°C (59°F). The lower supply water temperatures may produce improved results, but some data center operators will be concerned with managing condensation. Condensation is a valid concern and should be investigated with the installation of any direct cooling system.

A key advantage when using direct cooling is that adequate cooling is also possible when using much higher (e.g., 45°C [113°F]) supply water temperatures. In these situations, the direct-cooled components are still well below critical temperature as specified by the component manufacturer or server original equipment manufacturer.

We evaluated the thermal performance for a range of supply water temperatures in order to investigate the potential for cooling infrastructure energy savings using atypical cooling infrastructures; for example, dry coolers or cooling towers only. The temperature range used was 15°C–45°C (59°F–113°F).

In addition, because high supply water temperature can be used, the return water temperatures by definition are even higher—much higher than those of the return water in a typical data center. High return water temperatures open opportunities for heat reuse, providing additional energy savings (above that enabled by the energy reduction from the data center cooling infrastructure alone). The Asetek direct cooling system controls the return water temperature to take advantage of potential heat reuse opportunities by adjusting the water flow rate.

It should be noted that the return water temperature was an important parameter to vary, in order to understand the full range of thermal performance.

The remaining parameter varied was the server power level. The server power was adjusted by running High Performance LINPACK (HPL) software on the servers. Three power levels were used: Idle (no software applications running), 50 percent power (HPL operating at half the number of CPU cores), and Full or 100 percent power (HPL running all CPU cores).

The following parameter ranges were evaluated:

- Supply Water Temperature (15°C–45°C) [59°F–113°F]

- Water Flow Rate (~0.5 gallons per minute [gpm]–4.93 gpm) providing a return water temperature range of 17°C–52°C [63°F–126°F]

- Direct-Cooled IT Equipment Power

Idle	(120 watts/server)
50 Percent Power	(270 watts/server)
Full or 100 Percent Power	(430 watts/server)

See Appendix A for the parameter combinations tested.

2.4 Required Data

The key data required from Step 1 for the modeling in Step 2 were (1) measurements of electrical power (server cord power) for the base-case servers, and (2) liquid thermal heat transfer rates (total for liquid cooled servers) for the retrofit case. Server cord power relates to the electrical power flowing from a rack power distribution unit to the individual server via the power cord. The temperature of air entering the servers was also monitored, to provide a power correction for the air-cooled servers (details below). The data were collected at a frequency of 60

recordings per hour, providing a continuous and simultaneous time-series view of the parameters collected.

For this analysis, the fraction of server power cooled by the base case and direct-cooled case was required for the modeling software input. Electrical power was collected for all servers. For the base case of air-cooled servers, this electrical power was all that was needed because the fraction of cooling provided by the room air conditioning systems was 100 percent. The amount of water cooling for all direct cooled servers was recorded from a meter, allowing the fraction cooled by direct cooling to be calculated.

In addition, the front panel air-inlet temperature for each server was recorded. This information was used to adjust the energy use of the air-cooled servers if they were found to be significantly different than the average front panel air-inlet temperatures recorded for the direct cooled servers.

2.5 Test Setup

2.5.1 Physical Layout

The demonstration was held in a data center located on the Lawrence Berkeley National Laboratory campus in Building 50B, located in Berkeley, California. Figure 2-1 shows a view of the rack containing 38 Cisco servers equipped with the Asetek direct cooling technology. In the foreground is the Lytron water-to-water cooling distribution unit (CDU) used during the demonstration to adjust the supply water temperatures to simulate a variety of cooling infrastructures.

Figure 2-1: View Of Rack Containing 38 Cisco Servers Equipped With Asetek Direct Cooling Technology - Lytron CDU Shown In Foreground

Figure 2-2 shows a schematic diagram of the demonstration equipment layout. Thirty-eight Cisco servers were retrofitted with the technology being investigated. Two standard air-cooled Cisco servers, located in an adjacent rack, were used to provide baseline performance information. As mentioned, a Lytron water-to-water CDU was used to control the AsetekCDU supply water temperature. The Asetek equipment being evaluated consists of the RackCDU, flexible tubing going to each direct-cooled server, and the components inside each server providing component-level cooling.

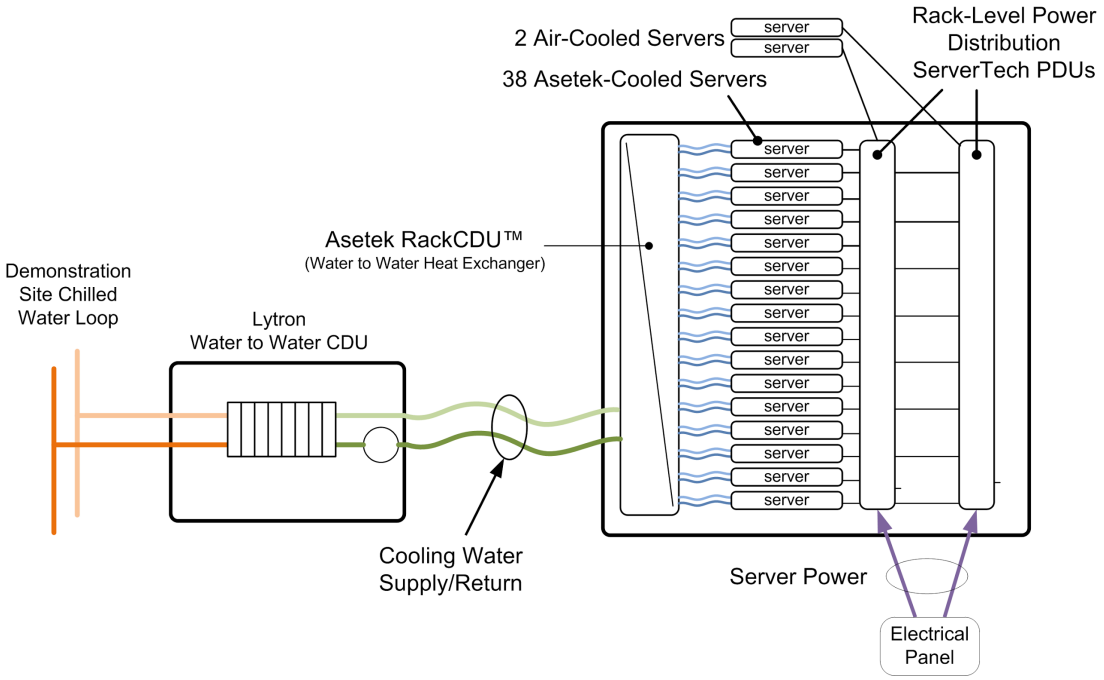

Figure 2-2: Demonstration Schematic Layout

2.5.2 Equipment Descriptions

Refer to Figure 2-3 as needed for the following described equipment.

IT Equipment (servers)

The IT equipment for this demonstration consisted of 40 1U servers. These servers were Cisco model UCS C220 M3 populated with two Intel Xeon 2690 v2 processors and 16 DIMMs: 8 DIMMs of 8GB Samsung M393B1K70DH0-YH9 and 8 DIMMs of 4GB Samsung M393B5270DH0-YK0. Thirty-eight of the servers were equipped with the Asetek cooling technology and installed in a rack containing the RackCDU. The Asetek technology was applied to the CPU and memory components. The two remaining servers, loaded in an adjacent rack, contained the same components but were not changed from the standard air-cooled model normally offered by Cisco. These two air-cooled servers were used to quantify the base case. Each Cisco server provided an intelligent platform management interface (IPMI) where the front panel air-inlet temperature, fan speeds, and CPU component temperatures were provided and recorded.

Lytron CDU

A Lytron CDU, model number LCS20, was used to supply variable water temperatures and flow rates to the RackCDU. The temperature set points were manually entered, and the water flow rate was manually controlled between both the building chilled water system and RackCDU as required for temperature stability and to adjust to the target RackCDU return temperature. A rear view of the Lytron CDU, with some piping details, is provided in Appendix B.

Asetek Direct Cooling System
The Asetek RackCDU D2C™ cooling system consisted of two parts: (1) a rack-mounted CDU providing cooling water distribution to the servers and water-to-water heat exchange between the server coolant and facilities water; and (2) cooling devices placed inside the server that contact temperature-sensitive components. For each server a set of flexible tubes provided a supply and return for cooling water going to the cooling devices inside the server.

Power Metering: Rack Power Distribution Units (ServerTech) and ION Meters
The power to each server was measured using a ServerTech power distribution unit (PDU) set consisting of model CWG-30VDE455 with an expansion PDU model CXG-30VDE455. These PDUs provided a simple network management protocol (SNMP) server used to gather data for the power consumption at each plug outlet. The total power consumed by the ServerTech PDUs were checked using data from two Schneider Electric PowerLogic ION Model 6200 power meters mounted on a nearby wall (Appendix C). Data from the ION and Onicon meters as shown in Figure 2-3 were collected using a Modbus serial network.

Btu Meters (Onicon System 10 Btu Meter)
There were two Onicon System 10 Btu meters. One meter between the Lytron CDU and RackCDU was set up for high-accuracy measurements and used a ¾-inch full-bore magnetic flow meter, model F-31AD-9A2 . The other Btu meter, located between the site chilled water loop and the Lytron CDU, used a standard model F-1300 turbine flow meter.

IT Equipment Thermal Stress Software
High-Performance LINPACK (HPL) version 2.0 benchmark software was used as a means of exercising the CPUs so as to vary the power consumed, and therefore the heat generated by the servers. The operating system used was Scientific Linux 6.4.

2.5.3 Data Collection

The values from a number of sensors were recorded during the demonstration. The time between records was one minute. Table 2-1 provides a listing of the variables recorded, with some additional comments. Figure 2-3 provides a simplified schematic showing the location of data points listed in Table 2-1, along with the data flow paths to the Joulex data collection software and database.

Name	Source	Purpose
Facility Water Supply Temperature	Btu Meter 1	backup reference
Facility Water Return Temperature	Btu Meter 1	backup reference
Facility Water Flow Rate	Btu Meter 1	monitor to assist with stability of the direct cooling supply water temperature
Facility Water Btu Rate	Btu Meter 1	monitor during testing to gauge temperature stability
RackCDU Water Supply Temperature	Btu Meter 2	record direct cooling supply water temperature
RackCDU Water Return Temperature	Btu Meter 2	record direct cooling return water temperature
RackCDU Water Flow Rate	Btu Meter 2	monitor direct cooling water flow rate
RackCDU Water Btu Rate	Btu Meter 2	record total direct cooling provided for direct-cooled servers
Individual Server Power Consumption Rate	Rack PDU 1 and 2	calculate direct-cooled and air-cooled server electrical power
Individual Server Front Panel Air Temperature	IPMI	adjust the air cooled server baseline power consumption values
Individual Server CPU 1 and 2 Temperature	IPMI	check for CPU component over temperature conditions
Individual Server Fan Speeds	IPMI	confirm fan speed BIOS settings
Rack PDU 1 and 2 Electrical Power	Power Meter 1 and 2	cross check the power reported by the rack PDUs

Table 2-1: Collected Data Descriptions

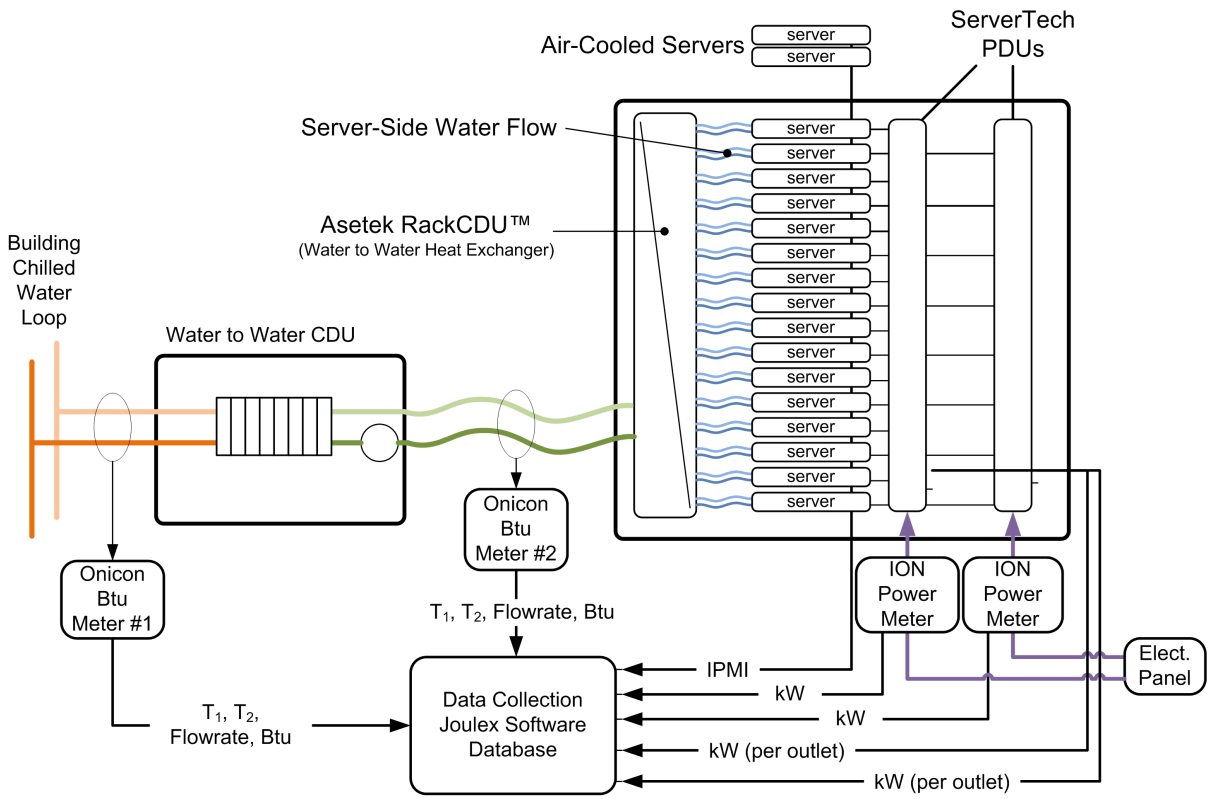

Figure 2-3: Schematic: Data Collection

The Joulex software data collection and storage system provided a means to collect each parameter at a frequency of once per minute. The Joulex database was later queried to obtain JavaScript Object Notation (JSON) files containing time-series record values for all parameters.

2.5.4 Test Procedure

The following outlines the general process used to adjust the test setup and gather the data for each set of parameters. Combinations of supply temperature, target return temperature and server power level were adjusted for each parameter set test. The combinations are listed in Appendix A

First, the server power level was set using the HPL software. For the idle power, level software applications and HPL processes were suspended. Second, the water supply temperature was set at the Lytron CDU control panel. Third, the water flow rate was manually adjusted until the desired return water temperature was reached. For every combination of server power and supply water temperature, a maximum flow rate test was completed. The maximum flow rate provided by the Lytron CDU was approximately 4.9 gpm. The Asetek system included control

electronics and a valve that modulated the water flow rate until the desired return water temperature was attained. This control system was not used during the demonstration.

After the thermal readings appeared to be steady, the setup was allowed to run for at least 30 minutes, thereby recording at least 30 sets of raw data.

The data center used for this demonstration was cooled using an underfloor pressurized cold-air plenum system fed by computer room air conditioning units (CRACs). This system had a tendency to supply cooler air to servers located near the bottom of a rack compared to those near the top. In an attempt to reduce this difference and improve the accuracy of the results, perforated floor tiles were replaced with solid floor tiles in the cold aisle, near the racks used for this demonstration.

This floor-tile change resulted in a more uniform and higher average server air inlet temperature recorded for the water-cooled servers compared to the first series of tests. The server air inlet temperature target was 27°C–29°C (81°F–84°F). This was at or slightly above the listing for class A1–A4 of the ASHRAE 2011 Thermal Guidelines, Table 4, recommended range. The average direct-cooled server inlet air temperatures were in a range of approximately 27°C–29°C (80°F–84°F), thereby achieving the ASHRAE target. The server air inlet temperatures at the air-cooled servers were lower and were used to calculate power consumption correction for the air-cooled servers to normalize power levels for the environmental conditions.

2.5.5 Calculation Process

The researcher team processed the thermal and electrical power measurements to determine the final results in approximately the following order:

1. Measured the heat energy flow captured by the direct cooling system
2. Measured electrical power consumed by the servers
3. Corrected the air-cooled server power measurements
4. Calculated the heat fraction removed by the direct cooling system
5. Developed data center energy models
6. Used the data center energy models to calculate the final results

Data were collected electronically at a frequency of 60 samples per hour and stored in the Joulex database.

The data were extracted from the database in a JSON file format and converted to a comma separated value (CSV) text file consisting of a row for each minute and a column for each data

parameter. Approximately 260 data parameters were collected; a large number of these parameters were values collected using IPMI from each server that were not used in the analysis.
A "test period" was determined by viewing the raw data and identifying a period, usually 20 to 30 minutes or more, when the thermal and power values appeared to have very little variation.

Heat Captured by Direct Cooling
The Btu rate was recorded directly from the Onicon meter #2, as shown in Figure 2-3.

Electrical Energy Consumed by the Direct-Cooled Servers
The power consumed by each direct-cooled server was recorded, for each outlet, from ServerTech PDU #1 and #2, as shown in Figure 2-3.

Electrical Energy Consumed by the Air-Cooled Servers
The power consumed by each air-cooled server (base case) was measured by the Server Technology, Inc. PDUs shown in Figure 2-3. The electrical power used in the final calculations involved adding a correction to these mean values, as described below.

Air-Cooled Server Electrical Power Correction
The air-cooled servers were located at the bottom of an adjacent rack and had a lower average air-inlet temperature compared to the average for the direct-cooled servers.

Server power increase as air inlet temperatures rise is often reported (Moss, Bean 2011). Information from this reference was used to make a correction (Appendix D). This correction resulted in a base case server power addition ranging from approximately 3 to 6 percent.

Note: The difference in server energy use between the base case and the retrofit case made a significant difference in the overall site energy savings, as the IT power consumed along with the cooling required were both affected. When considering a retrofit using this technology the change in IT equipment energy use should be verified.

Calculation of the Heat Fraction Removed by the Direct Cooling System
The fraction cooled by the direct cooling system for each period was determined by calculating the mean of the IT power measurements divided into the corresponding heat removed by the direct cooling system for each minute of a test period.

Determining the fraction of heat captured by direct cooling is necessary to evaluate the overall potential energy savings. The infrastructure energy savings comes from cooling this fraction with the direct-cooling system, and thereby reducing the cooling load on the existing and less-efficient room cooling system, as discussed later.

The standard error range for the percent of heat captured was calculated for each test period. These bands were small (2–4 percent). Therefore the mean values were used as input to the data

center energy use models. This standard error statistical process appears in Appendix E. The results of thermal testing appears in Appendix F, along with the standard error values for heat captured. The supply and return water temperatures used as part of the model inputs were the mean value of the measurements for each test period.

2.6 Data Center Energy Use Models

A key goal was to estimate the data center site level potential energy savings provided by this technology. A method was needed to convert the thermal performance results, using methods described above, into an estimate of the total energy consumed at the site level. This was accomplished by entering the test results into four different Romonet models.

It was assumed that the data center was cooled using CRAHs in the room supplied with cooling water coming from an infrastructure consisting of a chiller with a water-cooled condenser and cooling tower. This configuration was the base case model; the same cooling tower and chiller are used in the other three models, as described in the list below. The climate data selected for use in the model was for San Jose, California.

Supplying higher-temperature water to the data center for IT equipment cooling may have a number of energy-saving advantages, depending on the cooling plant. For example a plant consisting of cooling towers and chillers can usually supply a given amount of cooling more efficiently at a higher supply water temperature than it can at a lower water temperature.

The cooling tower case consisted of adding a dedicated cooling tower and chiller boost feature to the base case. The model used the chiller boost feature when necessary to maintain supply water temperature set point.

A much better scenario would be to use a dry cooler to supply what is needed for the direct cooling process. In many climates a dry cooler can be used for all hours during the year, to reject the heat from the server racks using this technology without using a chiller or cooling tower.

In addition, if high water temperatures are supplied to the device, even higher temperatures can be returned, possibly enabling heat reuse opportunities. This study did not explore the net economic gain or efficiencies associated with heat reuse, but the results of using water temperature combinations that might be used for these cases were measured, and are included along with the other results.

The Romonet modeling software was configured to analyze four data center configurations:

- Base Configuration
- Cooling Tower with Chiller Boost
- Dry Cooler with Chiller Boost

- Existing Chiller

Base Configuration is used to estimate the energy use for a data center equipped with air-cooled servers along with an unmodified cooling infrastructure.

The following are the model attributes for the base case shown in Figure 2-4 and the other three configurations:

Base Case IT power Consumption: 2 megawatts (MW) (Note: IT power consumption is different for the direct-cooled servers)
Room Air Conditioning: computer room air handlers (CRAHs) with an N + 20% capability and fixed fan speeds, CRAH units selectable depending on load.
Cooling Towers: N + 1 Counter-Flow
Chillers: N+1 Water Cooled
Uninterruptible Power Supply (UPS): N+1
Power Distribution Units: 2N Step-Down

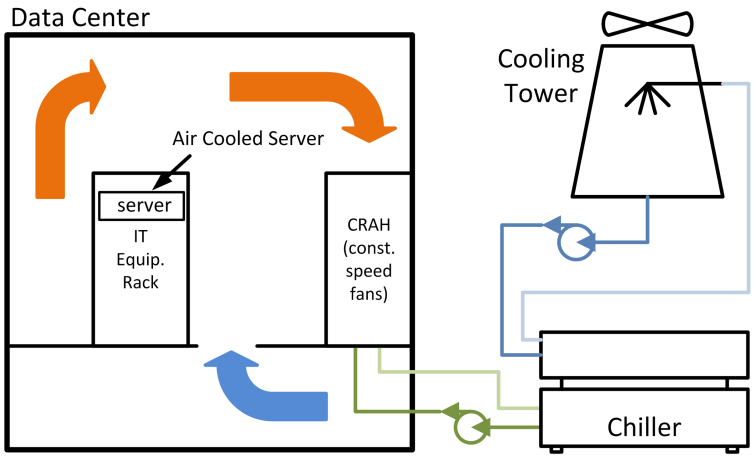

Figure 2-4: Base Case Configuration Cooling Infrastructure

Cooling Tower with Chiller Boost is used to estimate the energy use at a data center originally equipped with air-cooled servers. In this model, air-cooled servers are removed and replaced with direct-cooled servers, with the cooling water needed for the direct-cooled servers supplied by adding cooling towers. A liquid-to-liquid heat exchanger is added between the Rack CDUs and cooling tower water to provide the needed separation between the raw cooling tower water and water supplied for direct cooling. Shown in Figure 2-5 is a Chiller Boost liquid-to-liquid heat exchanger that provides additional cooling when climate conditions do not allow meeting the RackCDU set point. The Romonet model simulates using water from a chiller only when needed, which would be a small number of hours per year. This model assumes the existing cooling infrastructure stays in place and operational. It is not resized for the reduced maximum load.

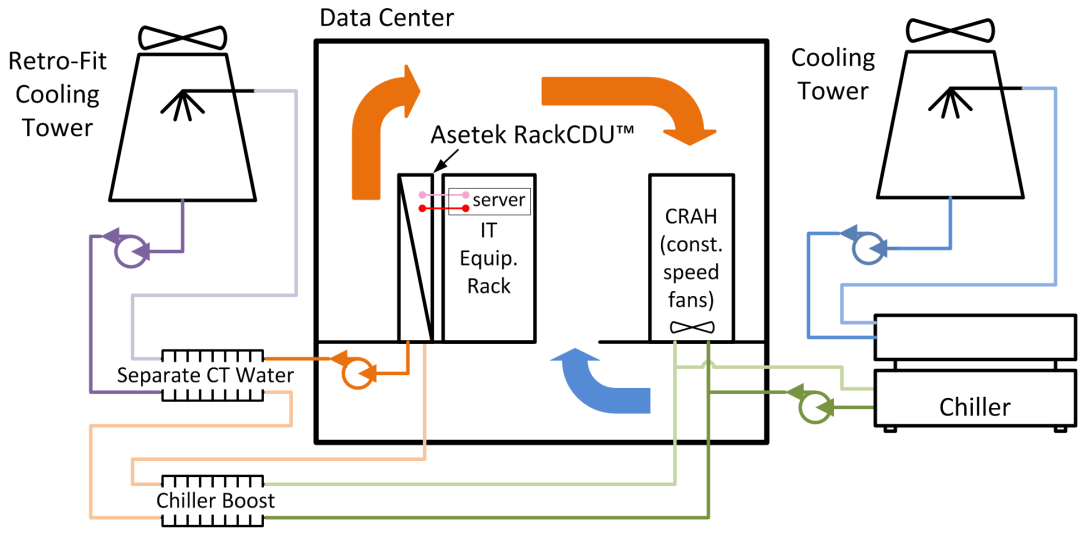

Figure 2-5: Cooling Tower and Cooling Tower With Chiller Boost Cooling Infrastructure Case Model

Dry Cooler with Chiller Boost is used to estimate the energy use at a data center originally equipped with air-cooled servers and base case infrastructure as described for the Cooling Tower case above, but the cooling tower is replaced by a dry cooler. Figure 2-6 also shows the Chiller Boost feature, used as described above. A dry cooler has the advantage of not requiring water. This model assumes the existing cooling infrastructure stays in place and operational. It is not resized for the reduced maximum load.

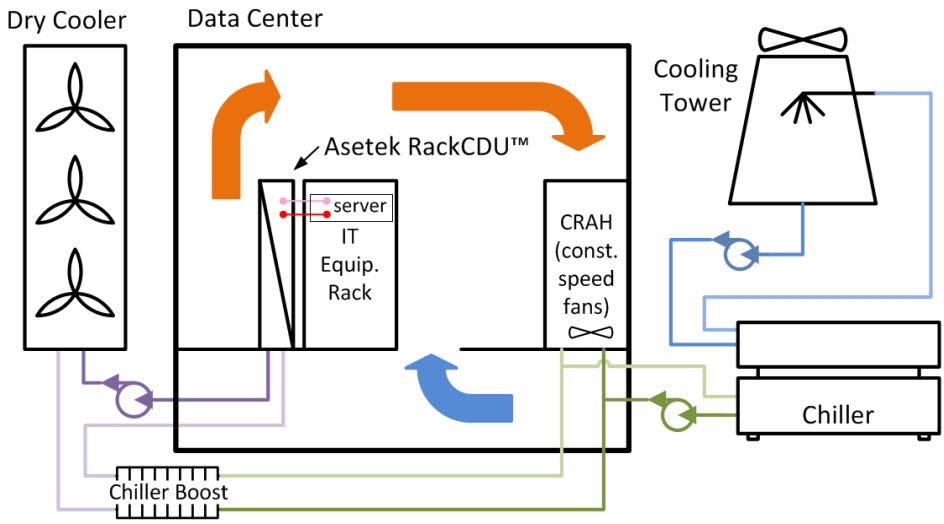

Figure 2-6: Dry Cooler And Dry Cooler With Chiller Boost Cooling Infrastructure Case Model

Existing Chiller Case is used to estimate the energy use for a data center originally equipped with air-cooled servers. The air-cooled servers are removed and replaced with direct-cooled servers. The cooling needed for the direct-cooled servers is supplied by connecting to the existing chilled water system (Figure 2-7).

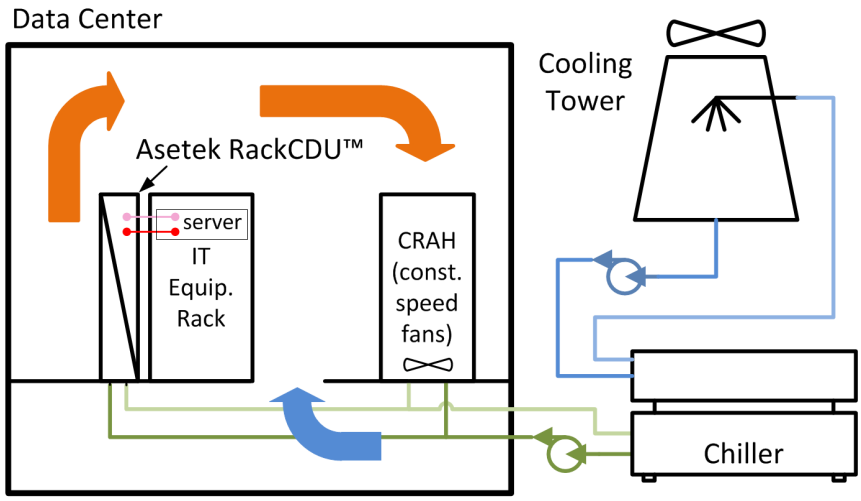

Figure 2-7: Direct Cooling Provided By Existing Chiller Infrastructure Case Model

Model Outputs

Using the four data center model configurations, the site-level energy consumption components and final results were calculated for the parameter combinations listed in Table 2-2. The RackCDU return water temperature targets were 10°C (50°F) above the supply target. The Asetek-equipped server heat-captured percentages used for site-level power modeling are in italicized rows (Appendix F).

Server Power Level	40°C Target RackCDU Supply Temperature	30°C Target RackCDU Supply Temperature	20°C Target RackCDU Supply Temperature	7°C Water Supply Temperature
100%	Dry Cooler	Cooling Tower, Dry Cooler w/Boost	Cooling Tower w/Boost, Dry Cooler w/Boost, Existing Chiller	Base Case CRAH
50%	Dry Cooler	Cooling Tower, Dry Cooler w/Boost	Cooling Tower w/Boost, Dry Cooler w/Boost, Existing Chiller	Base Case CRAH
Idle			Cooling Tower, Dry Cooler, Existing Chiller	Base Case CRAH

Table 2-2: Parameter Combinations Used For Site Energy Use Estimates

CHAPTER 3:
Results

The results of the thermal testing are presented in two sections: (1) thermal performance for maximum heat captured, and (2) thermal performance considering heat reuse applications. The overall data center energy use estimates for the three server power levels, using the thermal results for maximum heat captured, are then presented.

Appendix G presents the raw data for the three different server power levels used in the demonstration. These data were analyzed to determine heat capture fractions. Site energy consumption estimate models used the results for heat captured to estimate the overall site energy use for direct cooling compared to base cases.

The results are presented to address three questions:

- What are the thermal performance characteristics for maximum heat removed?
- What are the thermal performance characteristics considering heat reuse?
- How is the overall site energy affected if this technology replaces the air-cooled-only technology using performance for maximum cooling?

3.1 Maximum Heat Removed Using Direct Cooling

The results for maximum cooling provided by the Asetek technology are presented as a stand-alone attribute and then later used in site energy use estimates.

The prototype version tested had a constant water flow rate on the server side (Figure 2-3 "Server Side"). Colder water supply and/or higher flow rates of supply water on the facility side of the RackCDU (Figure 2-2 "cooling water supply/return") will extract the most heat (maximum cooling or heat captured) from the IT equipment for a given cooling water supply temperature and server load.

The results for the maximum facility-side flow rate (~4.93 gpm) for the facility side that could be provided by the demonstration setup are shown in Figure 3-1.

Server power level and supply water temperature were the parameters varied. The results are graphed in Figure 3-1, with the following observations:

- **For all power levels, the fraction of heat captured increases as the supply water temperature is lowered.**

This is expected, since surfaces of the cooling components (e.g., heat sinks, tubing, RackCDU) are colder relative to electronic component surfaces, air, or other structures that these components contact. Higher temperature differences cause higher heat transfer rates.

- **At higher supply water temperatures, the fraction of heat captured becomes distinctly different between the three power levels. At lower supply water temperatures, the fraction of heat captured is less different.**
  This suggests that the server power level should be considered when investigating the supply water temperature for an optimized overall operating cost.

  For example; if the server power level is low, and the supply water temperature is high (30°C; 86°F), the amount of heat captured is low compared to higher server power levels.

  It would be an advantage to be able to freely select from a wide range of supply water temperatures and capture a large fraction of the server power. However, the data show that this freedom is not available.

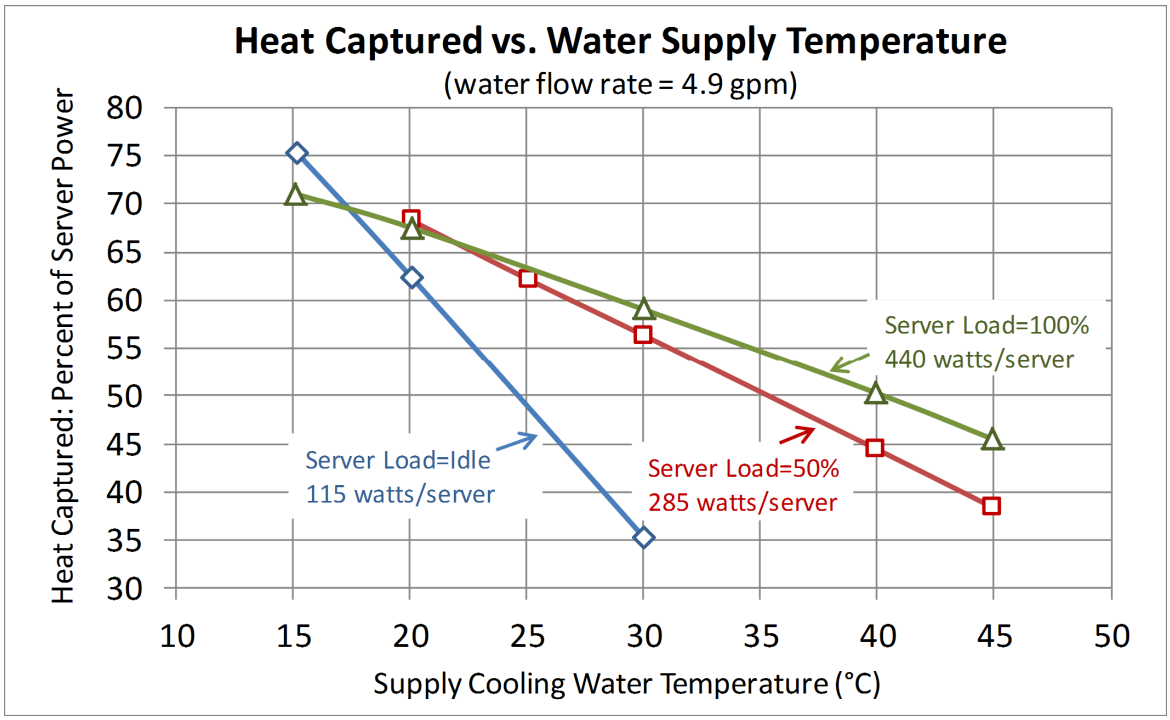

Figure 3-1: Maximum Heat Capture Vs. Water Supply Temperature For Three Sever Power Levels

3.2 Heat Reuse from Direct Cooling

The temperature of the water returned from the RackCDU was a key parameter for heat reuse. For a given supply water temperature, the return water temperature was adjusted by reducing the water flow rate until the desired return water temperature was reached.

The percent of heat captured for all power levels tested was plotted as a function of the return water temperature (Figure 3-2).

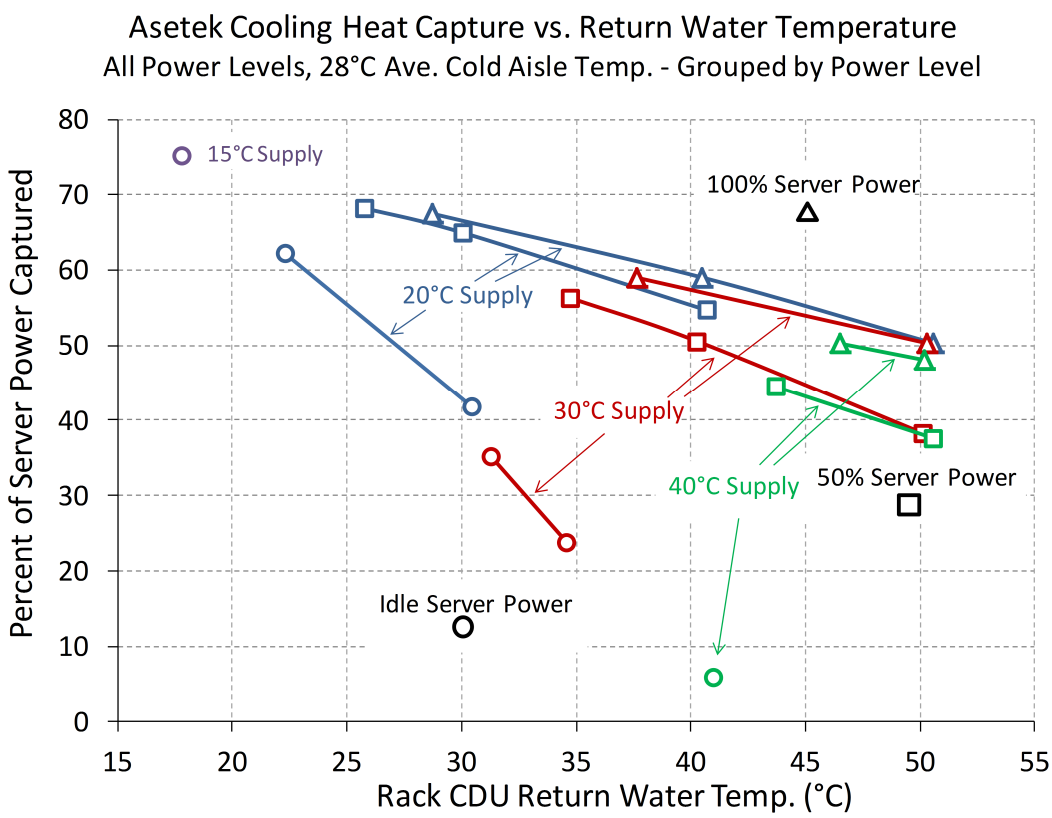

Figure 3-2: Heat Captured Vs. Return And Supply Water Temperature

In Figure 3-2 each line indicates the percent of server power captured as a function of return water temperature for a given server power level. The leftmost point on each supply temperature line corresponds to the maximum AsetekCDU facility side flow rate of 4.9 gallons per minute. Points going to the right on the same supply temperature line correspond to reduced flow rates. For each supply temperature line the flow rate is reduced and the return water temperature increases. The single point indicated for 40°C Supply (Idle Server Power) was

obtained with a 0.55 gpm water flow rate. Lowering the flow rate for this condition to obtain other points was not performed because this would result in an insignificant positive percentage of heat captured. The single point indicated for 15°C Supply (Idle Server Power) was obtained with a 4.9 gpm water flow rate.

Note: The server CPU and other component temperatures were within the internal temperature limits set by the server manufacturer for all results.

Observations:

- **For a given supply water temperature, a higher flow rate corresponds to a greater fraction of heat captured.** This behavior was expected because the water flow rate on the server side was constant, with the result that higher flow on the low-temperature side in turn dropped the temperature on the server side, thereby extracting a higher fraction of the heat from the direct-cooled components.

- **When considering reuse of the server heat, the return water temperature is the most important variable.** For given return water temperatures, the supply water temperature appears to have limited influence on the amount of heat captured. This result was not expected because it was assumed that the heat transfer rates would be strongly correlated with the temperature difference between the supply and return water, not just the return water temperature.

- **Setting a desired return water temperature for heat reuse applications results in an inconsistent heat capture percentage across the power levels.** Higher server power levels provide improved flexibility for supporting heat reuse applications compared to low server power levels.

3.3 Site Energy Consumption

The results of retrofit scenarios using the Asetek cooling technology were analyzed for each server power-level category (100 percent, 50 percent, and Idle) using software modeling. The model output details are presented in Appendix H.

The maximum cooling thermal performance results were used as input to the software models to estimate the site energy use differences comparing a base case to a number of retrofit cases, each using a different cooling infrastructure configuration.

Each server power-level result graphic (Figures 3-3 to 3-5) contains a number of stacked bar elements, with the base case on the left, dry cooler and cooling tower cases in the middle, and the existing chiller case on the right. The water temperature supplied to the RackCDU and configuration is indicated along the X axis. The estimate of the total annual site energy consumption is shown on the Y axis.

Site Energy Estimate: 100% Server Power

The overall energy reduction compared to the base case was similar, ranging from 17 percent to 23 percent, across the combinations of supply water temperature and infrastructure configurations. The following three graphics (Figures 3-3 to 3-5) present the overall data center energy use for three power levels when using the direct cooling system at the maximum heat capture configuration.

An interesting result is shown for the Existing Chiller case for 100 percent IT load (Figure 3-3).

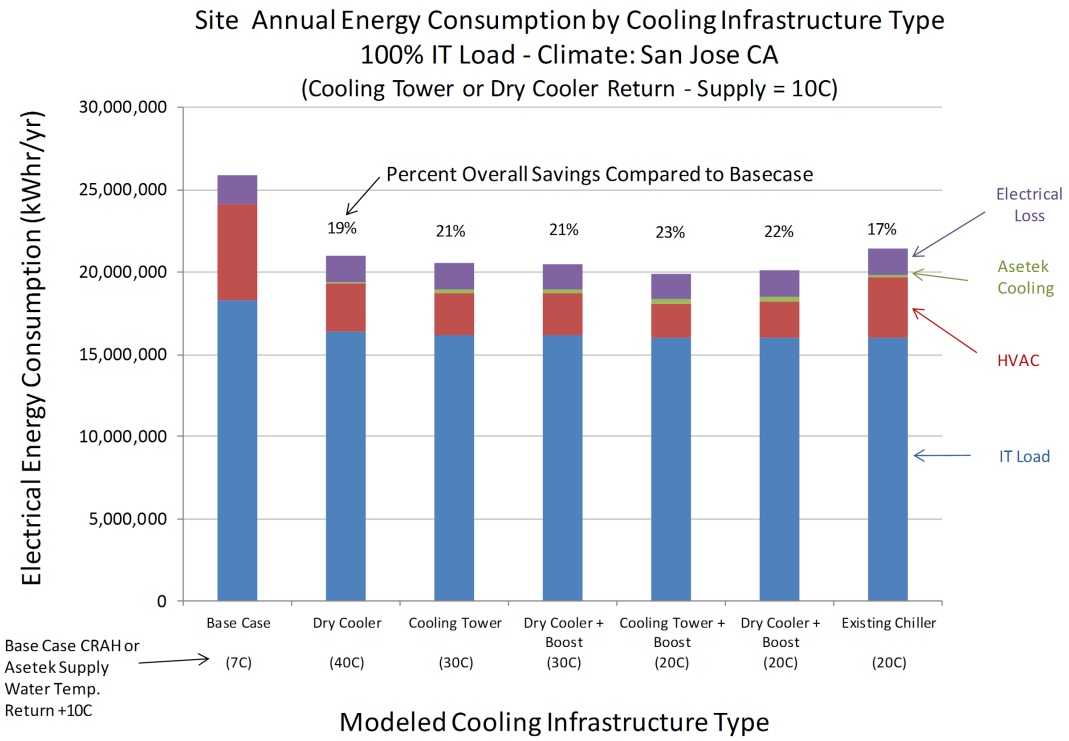

Figure 3-3: 100% Server Power Case Annual Energy Use Estimates

Because no new cooling plant components are required, this case is likely to provide a low first-cost retrofit compared to adding a dry cooler or cooling tower. The assumption is that a chilled water loop is available inside the data center and can be used for the direct cooling systems.

This lower first cost may result in a favorable total cost of ownership compared to adding plant components, offsetting slightly higher operating expenses.

Site Energy Estimate: 50% Server Power

Figure 3-4 shows the site energy savings results for tests using 50 percent server power. The overall energy saving results were slightly lower than the 100 percent server power results. This result was expected because the 100 percent and 50 percent heat captured was similar for the 20°C and 30°C (68°F and 86°F) supply water temperature cases (Figure 3-1).

It is interesting to note that 40 percent of the savings is due to the reduction in server power allowed by the direct cooling system.

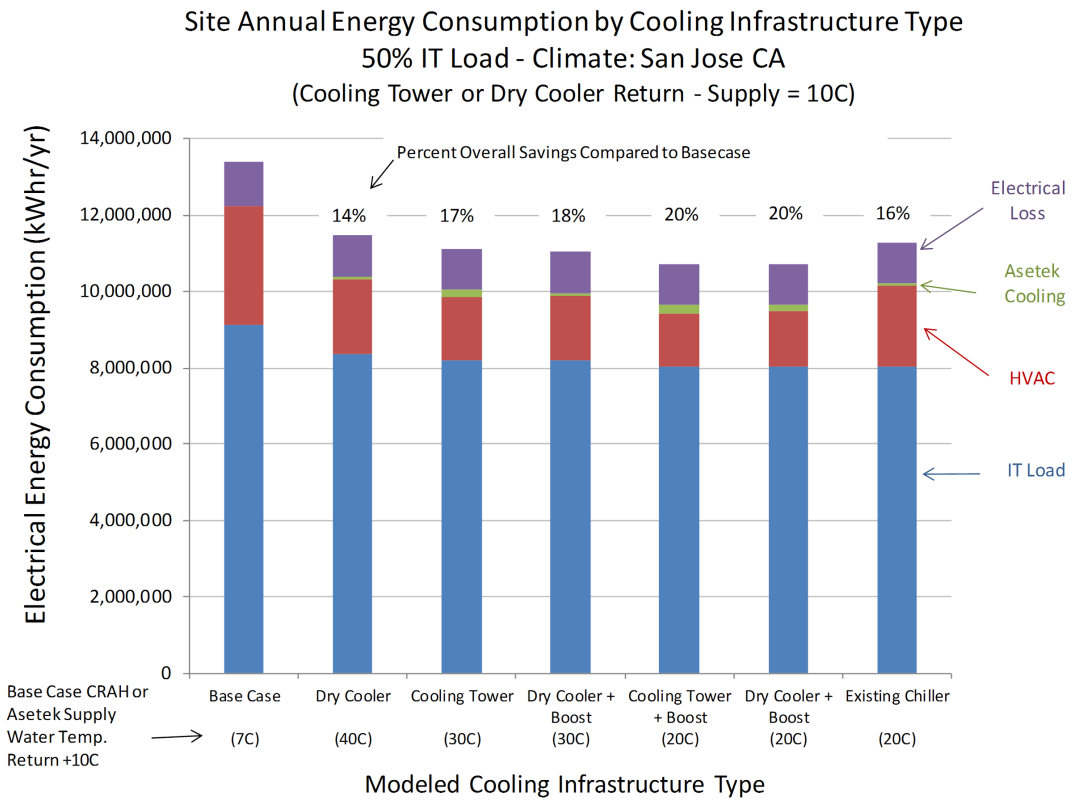

Figure 3-4: 50% Server Power Case Annual Energy Use Estimates

Site Energy Estimate: Idle Server Power

The site energy use results were distinctly different for the tests using the idle server power level (Figure 3-5). The heat captured using 20°C (68°F) supply water temperature was somewhat lower, so similar overall results were expected. But since the power distribution and cooling infrastructure losses were not reduced in the same ratio as the

server power reduction, the site-level savings were approximately half of that for the 100 percent and 50 percent server power cases.

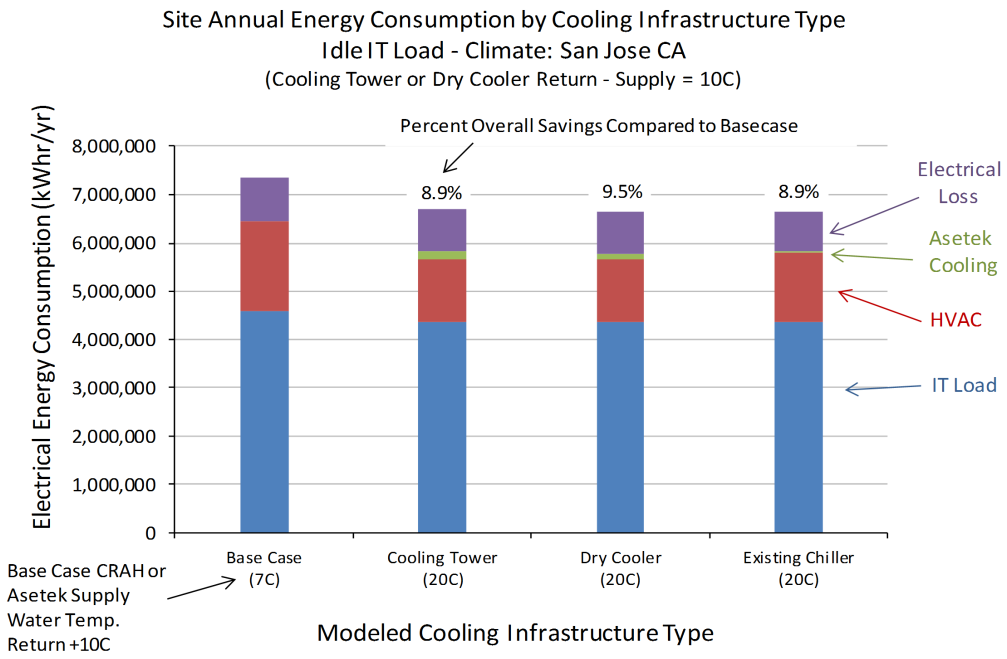

Figure 3-5: Idle Server Power Annual Energy Use Estimates

CHAPTER 4:
Conclusions and Recommendations

A methodology used to measure and quantify the heat captured by the Asetek cooling technology has been documented. The overall site energy consumption has been estimated using a retrofit consisting of this direct cooling technology and several variations of cooling plant infrastructure. This direct cooling technology should provide a significant reduction (~14–20 percent) in total data center site energy consumed if implemented as modeled for an average IT server load of 50 percent.

The existing chiller implementation, even with a reduced operational savings, may provide a favorable total cost of ownership compared to a retrofit that includes adding a cooling tower or dry cooler.

Applying this direct cooling technology to heat reuse applications should be analyzed closely. If the server power level is low, the heat that can be supplied may not justify the resources needed to provide that heat.

It is interesting to note that the reduction in server energy consumed is about 40 percent of the overall savings when compared to the base case. Therefore, as part of planning a retrofit using this technology, obtaining an accurate estimate of this in-server energy reduction, using specific hardware and software, is strongly recommended.

Capturing heat appears to be more difficult as the supply or return water temperature is increased. This is especially true for the case of high return water temperature.

Significant site-level energy savings (16–20 percent for the 50 percent server power case) were shown as feasible using this technology, even though it does not capture 100 percent of the server heat. Increasing the fraction of heat captured is desirable, especially if considering heat-reuse applications requiring high return water temperatures.

Understanding the heat paths for the server power not captured by the Asetek cooling system is an area suggested for future study.

A closed-form solution, based on measurements, that relates supply water temperature, return water temperature, ambient room conditions, and server power level to the heat captured would be very valuable for comparing different direct-cooling solutions similar to this technology. This effort is also an area recommended for future studies.

GLOSSARY

Term	Definition
Btu	British thermal unit
BIOS	basic input/output system
CDU	cooling distribution unit
CPU	central processing unit
CRAH	computer room air handler
DIMM	dual inline memory module
gpm	gallons per minute
IPMI	intelligent platform management interface
IT	information technology
kW	kilowatt
LBNL	Lawrence Berkeley National Laboratory
LLNL	Lawrence Livermore National Laboratory
PDU	power distribution unit
pPUE	partial power usage effectiveness
PUE	power usage effectiveness
ROI	return on investment
TCO	total cost of ownership
UPS	uninterruptible power supply
VFD	variable frequency drive
W	watt

REFERENCES

Koomey, J. 2011. *My new study of data center electricity use in 2010.* July 31. Retrieved February 7, 2014 from http://www.koomey.com/post/8323374335. (last accessed 4/18/2014)

Moss, David (Dell Inc.), and John Bean, Jr. (APC). 2011. Energy Impact of Increased Server Inlet Temperature, White Paper #138, http://www.apcdistributors.com/white-papers/Cooling/WP-138%20Energy%20Impact%20of%20Increased%20Server%20Inlet%20Temperature.pdf

Coles, Henry, and Steve Greenberg (Lawrence Berkeley National Laboratory), and Corinne Vita (Vigilent). 2012. *Demonstration of Intelligent Control and Fan Improvements in Computer Room Air Handlers*. Publication number LBNL-6007E.

APPENDIX A:
Test Plan Parameter Combinations

Table A-1: Test Plan Parameter Combinations For Servers At Idle Power Level

IT Power Level	Supply Water Temperature (°C)	Return Water Temperature (°C)	Flow Rate (gpm)
Idle	15	TBD	Full Flow
Idle	20	TBD	Full Flow
Idle	20	30	As Required
Idle	30	40	As Required
Idle	40	50	As Required

Table A-2: Test Plan Parameter Combinations For Servers At 50% Power Level

IT Power Level	Supply Water Temperature (°C)	Return Water Temperature (°C)	Flow Rate (gpm)
50%	20	TBD	Full Flow
50%	20	30	As Required
50%	20	40	As Required
50%	25	TBD	Full Flow
50%	30	TBD	Full Flow
50%	30	40	As Required
50%	30	50	As Required
50%	40	TBD	Full Flow
50%	40	50	As Required
50%	45	TBD	Full Flow

Table A-3: Test Plan Parameter Combinations For Servers At 100% Power Level

IT Power Level	Supply Water Temperature (°C)	Return Water Temperature (°C)	Flow Rate (gpm)
Full	15	TBD	Full Flow
Full	15	30	As Required
Full	15	40	As Required
Full	15	50	As Required
Full	20	TBD	Full Flow
Full	20	40	As Required
Full	20	50	As Required
Full	30	TBD	Full Flow
Full	30	50	As Required
Full	40	TBD	Full Flow
Full	40	50	As Required
Full	45	TBD	Full Flow

APPENDIX B:
Lytron CDU Plumbing Detail

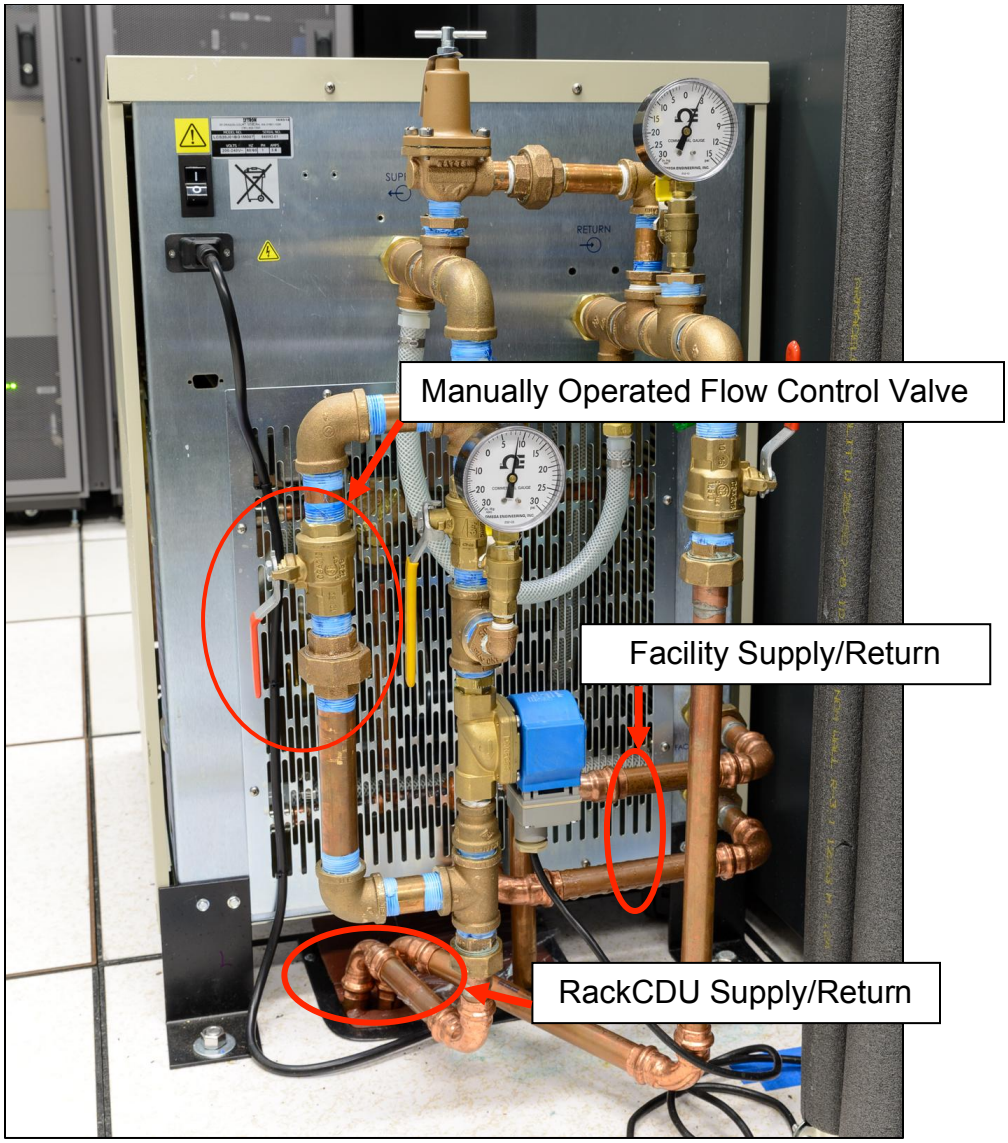

Figure B-1: Lytron CDU Piping Details

APPENDIX C:
Electrical Power and Btu Meters

Figure C-1: Installed Electrical And Btu Metering

APPENDIX D:
Air-Server Power Correction

Figure 5

Composite server power vs. inlet temperature

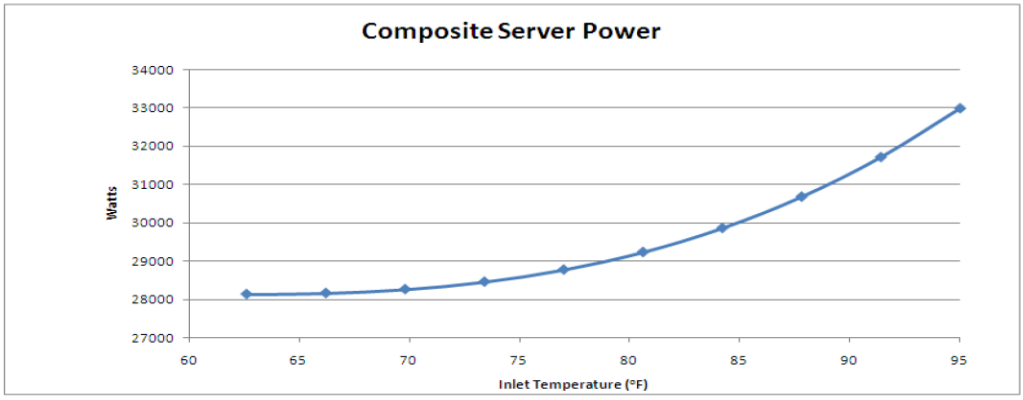

(Source: Moss, D., and J. Bean, Jr., White Paper 2011)

Figure D-1: Composite Server Power

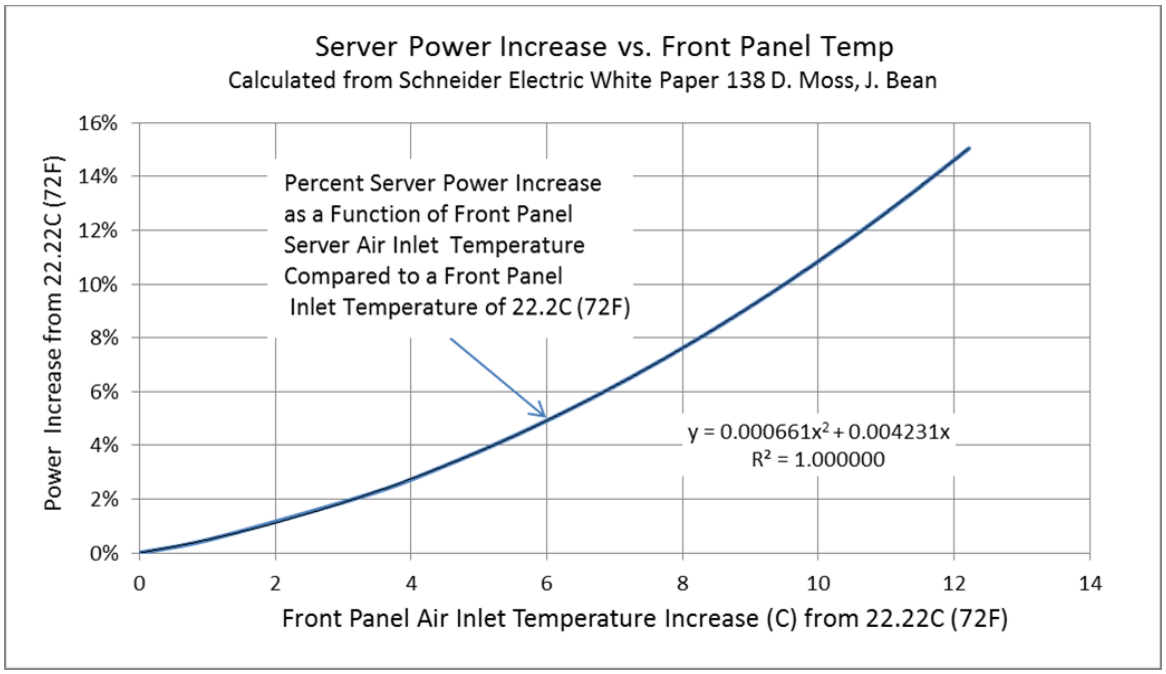

Figure D-2: Air-Server Power Front Panel Air Inlet Temperature Correction Graph, Using Data from Figure D-1

APPENDIX E:
Statistical Analysis of Fraction Cooled

The heat removed by the direct-cooling system for each set of parameters was determined by reading Btu meter #2 (refer to Figure 2-3) in one-minute intervals over the test period. The fraction cooled for each one-minute period was calculated using the reading from the Btu meter and dividing by the average power consumed by the direct-cooled servers. The readings involved were sequential, so an autocorrelation function was used to correct the quantity of readings to enable a calculation of the standard error. The calculation steps were:

1. Calculate the mean and standard deviation of the calculated fraction cooled values for the test period.

2. Use the R acf function to obtain nlag.

R code example:

ac <- acf(Full_45_50$Water_Cooling_Over_Total_Water_Server_Power)
The argument "Full_45_50$Water_Cooling_Over_Total_Water_Server_Power contains all the measurements.

3. Calculate nlag.
nlag = which (abs(ac$acf)<0.2)[1]

4. Calculate the effective number of measurements
Effective number = actual number / nlag

5. Calculate the Standard Error.
Standard Error = standard deviation of data points measured in test period / square root of the effective number of data points

6. Calculate the Upper and Lower Standard Error Limits.
The upper and lower limits of the Standard Error is defined as:
Upper 95% Limit = mean of measurements in the test period + (standard error x 1.96)
Lower 95% Limit = mean of measurements in the test period - (standard error x 1.96)

APPENDIX F:
Thermal Test Results

Rows italicized in blue were used as input for site energy use modeling.

Power Level	Mean Supply Temp. (C)	Mean Return Temp. (C)	Mean Water Flow Rate (gpm)	Fraction Cooled Directly (%)	Mean Power Water Servers (watts)	Mean Power Air Servers (watts)	Mean Water Servers Air Inlet (C)	Mean Air Servers Air Inlet (C)	Heat Captured Standard Error Conf. Band (%)
Idle	15.2	17.8	4.93	75.3	117	110	28.9	22.6	0.94
Idle	20.2	22.3	4.93	62.2	117	109	28.9	22.5	2.24
Idle	*20.0*	*30.5*	*0.68*	*41.7*	*117*	*110*	*28.6*	*22.0*	*4.38*
Idle	30.1	31.3	4.93	35.1	118	111	29.4	22.7	1.81
Idle	28.8	34.6	0.928	23.8	118	110	29.5	23.0	5.20
Idle	39.2	41.0	0.55277	5.8	119	109	29.4	22.5	0.39
0.5	20.1	40.7	1.09	54.7	281	289	29.2	22.3	5.17
0.5	*20.1*	*30.1*	*2.6*	*65.0*	*276*	*287*	*29.1*	*22.6*	*3.80*
0.5	20.1	25.8	4.93	68.2	277	290	28.5	22.5	1.48
0.5	25.1	30.3	4.93	62.1	279	290	28.1	22.4	0.81
0.5	29.8	50.1	0.796	38.3	288	289	28.1	22.5	3.14
0.5	*30.0*	*40.3*	*2.02*	*50.6*	*282*	*290*	*28.1*	*22.5*	*1.22*
0.5	30.1	34.8	4.93	56.3	281	289	28.1	22.4	2.72
0.5	*39.7*	*50.6*	*1.46*	*37.6*	*288*	*290*	*28.1*	*22.5*	*1.46*
0.5	40.0	43.7	4.91	44.5	285	289	28.1	22.5	0.91
0.5	44.9	48.2	4.93	38.4	287	290	28.3	22.8	1.80
1.0	15.2	24.2	4.93	71.0	428	445	28.7	23.8	1.02
1.0	15.2	30.1	2.76	67.4	417	443	28.8	23.2	0.96
1.0	15.2	40.4	1.45	59.2	422	447	28.8	23.3	1.18
1.0	15.3	50.0	0.902	49.9	427	448	28.5	23.5	8.08
1.0	20.1	50.6	1.05	50.3	434	440	29.4	24.0	2.06
1.0	*20.1*	*28.7*	*4.93*	*67.5*	*429*	*450*	*29.6*	*24.8*	*1.65*
1.0	20.1	40.5	1.84	59.0	432	445	29.6	24.7	1.31
1.0	29.9	50.3	1.58	50.3	436	448	29.4	24.0	1.38
1.0	*30.1*	*37.6*	*4.93*	*59.0*	*431*	*444*	*28.9*	*24.0*	*1.81*
1.0	39.9	50.2	3.01	48.1	438	447	28.8	23.9	1.40
1.0	*40.0*	*46.5*	*4.92*	*50.3*	*437*	*445*	*28.7*	*23.9*	*1.25*
1.0	44.9	50.9	4.9	45.5	438	441	28.4	23.6	1.16

Table F-1: Thermal Performance Test Results

APPENDIX G:
Data at 100%, 50%, and Idle Server Power

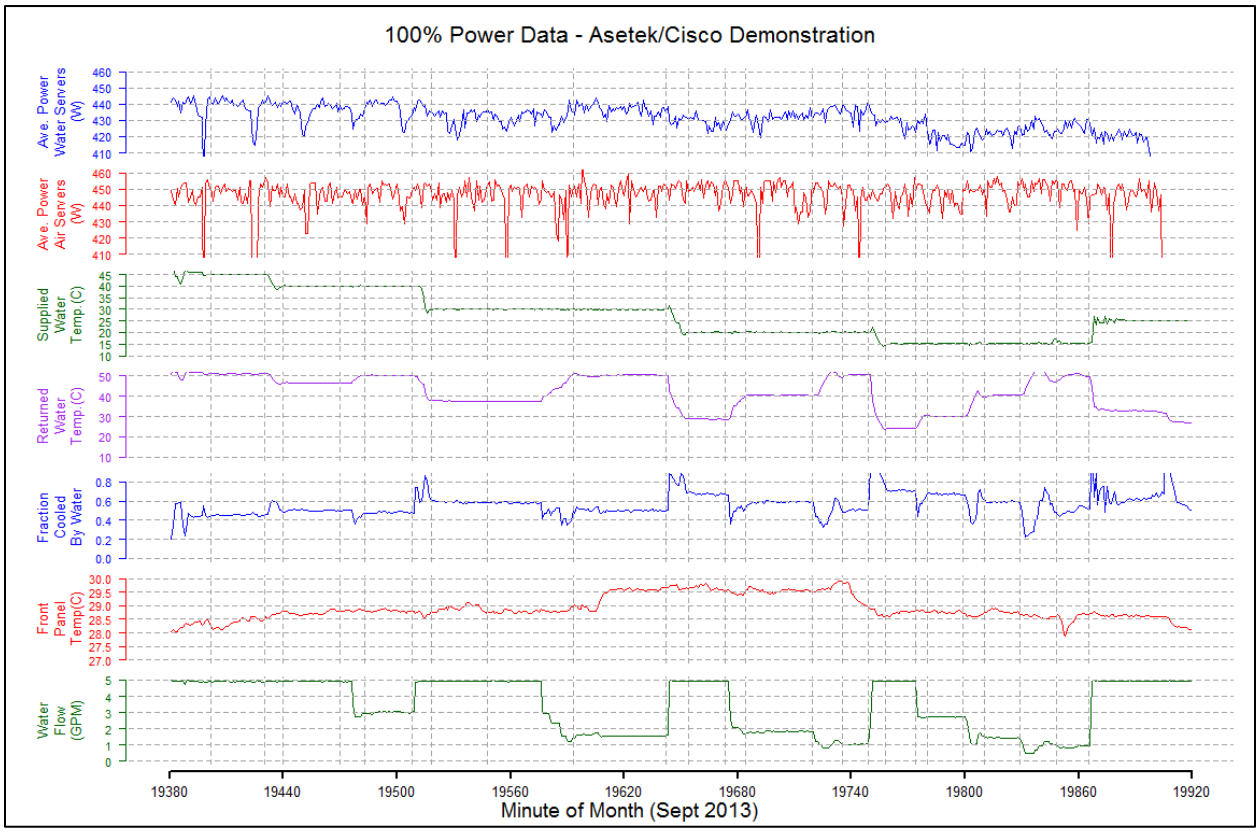

Figure G-1: Time Series Line Graphs - 100% Server Power Test Data

Data at 50 Percent Server Power

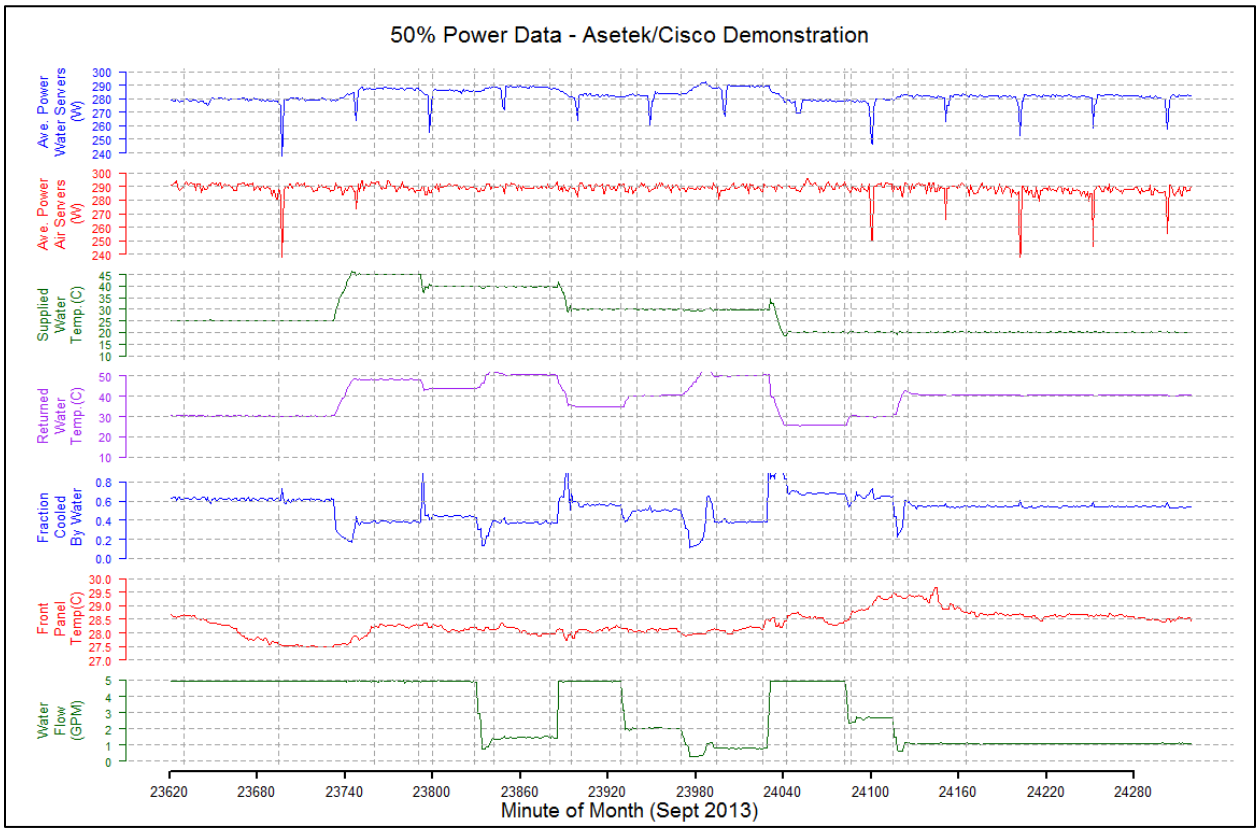

Figure G-2: Time Series Line Graphs - 50% Server Power Test Data

Data at Idle Server Power

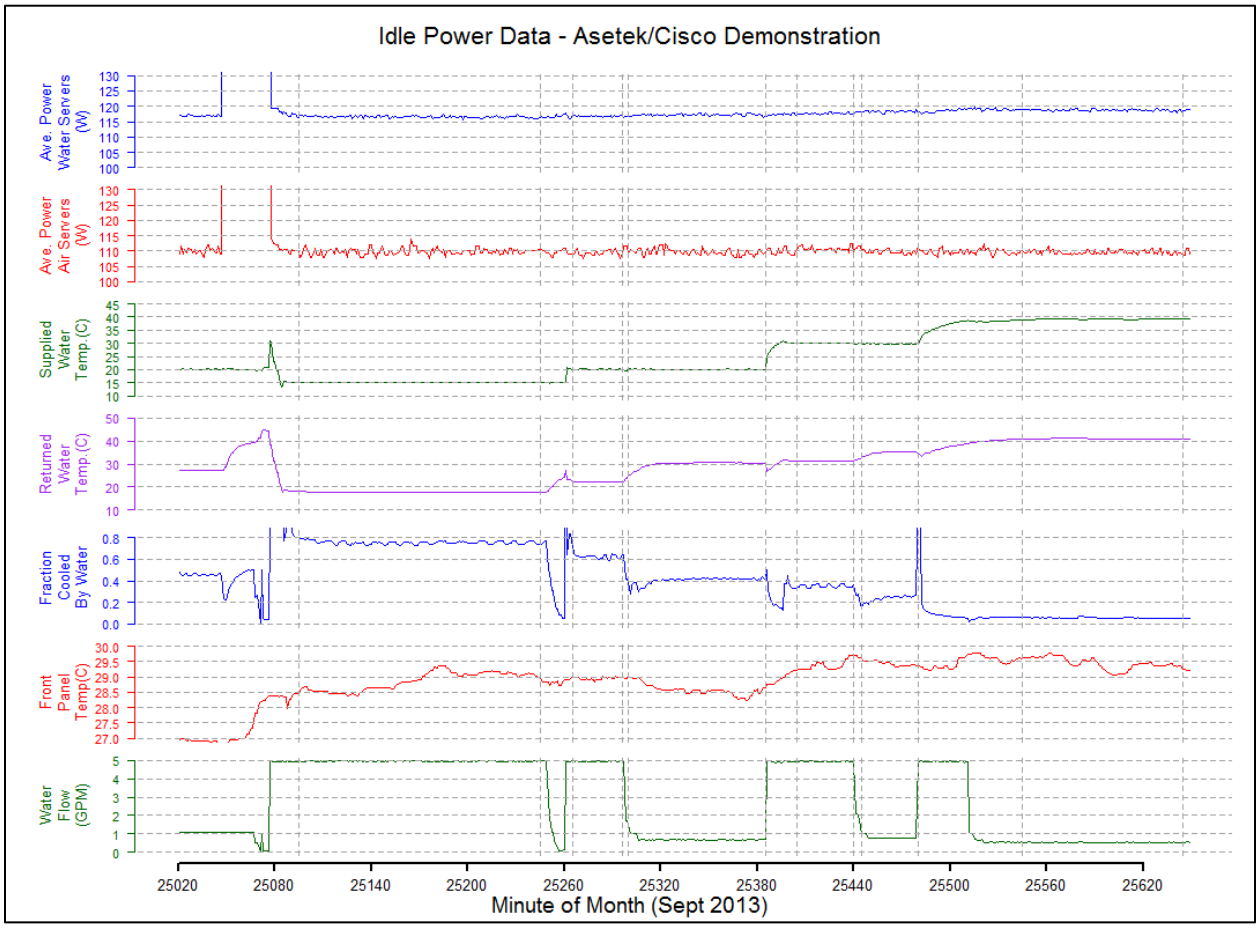

Figure G-3: Time Series Line Graphs - Idle Server Power Test Data

APPENDIX H:
Modeling Results: 100%, 50% and Idle Server Power

	100% IT Load						
	IT Load	HVAC	Liquid Cooling	Elect. Loss	Total	PUE	Fraction Captured
Base Case	1.83E+07	5.81E+06	0.00E+00	1.78E+06	2.59E+07	1.42	0.00
Dry Cooler 40°C Supply	1.64E+07	2.87E+06	1.07E+05	1.61E+06	2.10E+07	1.28	0.48
Cooling Tower 30°C Supply	1.62E+07	2.53E+06	2.76E+05	1.59E+06	2.06E+07	1.27	0.59
Dry Cooler w/Boost 30°C Supply	1.62E+07	2.54E+06	2.20E+05	1.59E+06	2.05E+07	1.27	0.59
Cooling Tower w/Boost 20°C Supply	1.61E+07	1.99E+06	2.83E+05	1.58E+06	1.99E+07	1.24	0.67
Dry Cooler w/Boost 20°C Supply	1.61E+07	2.16E+06	3.16E+05	1.58E+06	2.01E+07	1.25	0.67
Existing Chiller	1.61E+07	3.66E+06	1.08E+05	1.59E+06	2.14E+07	1.33	0.67

Table H-1: Model Results For 100% Server Power

Modeling Results: 50% Server Power

	50% IT Load						
	IT Load	HVAC	Liquid Cooling	Elect. Loss	Total	PUE	Fraction Captured
Base Case	9.14E+06	3.10E+06	0.00E+00	1.14E+06	1.34E+07	1.46	0.00
Dry Cooler 40°C Supply	8.37E+06	1.93E+06	7.95E+04	1.09E+06	1.15E+07	1.37	0.38
Cooling Tower 30°C Supply	8.20E+06	1.66E+06	1.80E+05	1.08E+06	1.11E+07	1.36	0.51
Dry Cooler w/Boost 30°C Supply	8.20E+06	1.67E+06	8.26E+04	1.08E+06	1.10E+07	1.35	0.51
Cooling Tower w/Boost 20C Supply	8.02E+06	1.41E+06	2.10E+05	1.07E+06	1.07E+07	1.33	0.65
Dry Cooler w/Boost 20°C Supply	8.02E+06	1.47E+06	1.53E+05	1.07E+06	1.07E+07	1.33	0.65
Existing Chiller	8.02E+06	2.13E+06	5.01E+04	1.07E+06	1.13E+07	1.41	0.65

Table H-2: Model Results For 50% Server Power

Modeling Results: Idle Server Power

	Idle IT Load						
	IT Load	HVAC	Liquid Cooling	Elect. Loss	Total	PUE	Fraction Captured
Base Case	4.59E+06	1.87E+06	0.00E+00	8.86E+05	7.35E+06	1.60	0.00
Cooling Tower w/Boost (20°C) Supply	4.38E+06	1.28E+06	1.58E+05	8.74E+05	6.69E+06	1.53	0.42
Dry Cooler w/Boost (20°C) Supply	4.38E+06	1.29E+06	1.09E+05	8.74E+05	6.65E+06	1.52	0.42
Existing Chiller	4.38E+06	1.42E+06	1.50E+04	8.74E+05	6.69E+06	1.53	0.42

Table H-3: Model Results For Idle Server Power

DISCLAIMER

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