



# Energy Performance Evaluation of Aquila's Aquarius Fixed Cold Plate Cooling System at NREL's High Performance Computing Center

David Sickinger,<sup>1</sup> David Martinez,<sup>2</sup> and Bob Bolz<sup>3</sup>

<sup>1</sup> *National Renewable Energy Laboratory*

<sup>2</sup> *Sandia National Laboratories*

<sup>3</sup> *Aquila*

**NREL is a national laboratory of the U.S. Department of Energy  
Office of Energy Efficiency & Renewable Energy  
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at [www.nrel.gov/publications](http://www.nrel.gov/publications).

Contract No. DE-AC36-08GO28308

**Technical Report**  
NREL/TP-2C00-73356  
February 2019



# Energy Performance Evaluation of Aquila's Aquarius Fixed Cold Plate Cooling System at NREL's High Performance Computing Center

David Sickinger,<sup>1</sup> David Martinez,<sup>2</sup> and Bob Bolz<sup>3</sup>

<sup>1</sup> *National Renewable Energy Laboratory*

<sup>2</sup> *Sandia National Laboratories*

<sup>3</sup> *Aquila*

## Suggested Citation

Sickinger, David, David Martinez, and Bob Bolz. 2019. *Energy Performance Evaluation of Aquila's Aquarius Fixed Cold Plate Cooling System at NREL's High Performance Computing Center*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-2C00-73356. <https://www.nrel.gov/docs/fy19osti/73356.pdf>.

**NREL is a national laboratory of the U.S. Department of Energy  
Office of Energy Efficiency & Renewable Energy  
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at [www.nrel.gov/publications](http://www.nrel.gov/publications).

Contract No. DE-AC36-08GO28308

**Technical Report**  
NREL/TP-2C00-73356  
February 2019

National Renewable Energy Laboratory  
15013 Denver West Parkway  
Golden, CO 80401  
303-275-3000 • [www.nrel.gov](http://www.nrel.gov)

## NOTICE

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. A portion of the research was performed using computational resources sponsored by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at [www.nrel.gov/publications](http://www.nrel.gov/publications).

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via [www.OSTI.gov](http://www.OSTI.gov).

*Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.*

NREL prints on paper that contains recycled content.

## Acknowledgments

Sandia National Laboratories made this project possible by providing the system hardware. Staff from Sandia involved in the setup included Jay Livesay, Stephen Monk, and David Smith.

Andree Jacobson, Froilan Ramirez, and Ryan Schecker from Aquila provided on-site support during the setup and configuration of the system.

Staff from the National Renewable Energy Laboratory involved in the setup and test runs of the system included Surendra Sunkari and Roy Fraley. We also thank Otto Van Geet, who provided valuable input and constructive feedback along the way.

## Nomenclature and List of Acronyms

BTU	British thermal unit
CDU	Cooling distribution unit
DIMM	Dual in-line memory module
ERW	Energy recovery water
ESIF	Energy Systems Integration Facility
HPC	High-performance computing
HSPB	Hot Swap Power Board
I <sup>2</sup> C	Inter-Integrated Circuit
IMPI	Intelligent Management Platform Interface
IT	Information technology
kW	kilowatt (1,000 watts)
LAMMPS	Large-scale Atomic/Molecular Massively Parallel Simulator
NREL	National Renewable Energy Laboratory
OCP	Open Compute Project
PDU	Power distribution unit
PUE	Power usage effectiveness
RAM	Random-access memory
SSD	Solid-state drive
Teraflop	10 to the 12th power floating point calculations per second
TIM	Thermal interface material
VAC	Volts, alternating current
VDC	Volts, direct current
W	watt

## Executive Summary

In the first half of 2018, as part of a partnership with Sandia National Laboratories, Aquila installed its fixed cold plate, liquid-cooled Aquarius rack solution for high-performance computing (HPC) clustering at the National Renewable Energy Laboratory's (NREL's) Energy Systems Integration Facility (ESIF). This new fixed cold plate, warm-water cooling technology together with a manifold design provides easy access to service nodes and eliminates the need for server auxiliary fans altogether. Aquila and Sandia National Laboratories chose NREL's HPC Data Center for the initial installation and evaluation because the data center is configured for liquid cooling and has the required instrumentation to measure flow and temperature differences to facilitate testing. This paper gives an overview of the Aquarius fixed cold plate cooling technology and provides results from early energy performance evaluation testing.

Sandia's Aquila-based HPC cluster was named "Yacumama" and was configured to operate independently from all other HPC systems in the ESIF data center. There are 36 compute nodes that have INTEL S2600KP motherboards. The motherboards are configured with dual X86\_64 XEON central processing units, 128 GB of random access memory (RAM), 128 GB solid-state drive (SSD), and an Omni-Path adapter. The supplied configuration is capable of providing >40 teraflops of LINPACK performance while drawing less than 15 kW of power.

In building the data center, NREL's vision was to create a showcase facility that demonstrates best practices in data center sustainability and serves as an exemplar for the community. The innovation was realized by adopting a holistic "chips to bricks" approach to the data center, focusing on three critical aspects of data center sustainability: (1) efficiently cool the information technology equipment using direct, component-level liquid cooling with a power usage effectiveness design target of 1.06 or better; (2) capture and reuse the waste heat produced; and (3) minimize the water used as part of the cooling process. There is no compressor-based cooling system for NREL's HPC data center. Cooling liquid is supplied indirectly from cooling towers.

The Yacumama cluster installation was straightforward and easily integrated directly into the data center's existing hydronic system. A round of discovery testing was conducted to identify the range of reasonable supply temperatures to the fixed cold plates and the impact of adjusting facility flow. Then LINPACK tests at 100% duty cycle were run for 48 hours. Results are provided in Section 3, and the key takeaway is that this fixed cold plate design provides a very high percentage of heat capture direct to water—up to 98.3% when evaluating compute nodes only (the percentage drops to 93.4% when evaluating compute nodes along with the Powershell for the system because the Powershell is not direct liquid cooled).

This cluster has been in operation for nearly 10 months, requiring zero maintenance, and no water leaks were observed. The Yacumama system will be returned to service at Sandia's recently completed warm-water-cooled HPC data center in early 2019.

# Table of Contents

<b>1</b>	<b>Introduction</b> .....	<b>1</b>
1.1	Brief Aquarius Development History.....	1
1.2	Aquila’s Aquarius Cooling System.....	2
1.3	Rack Hardware.....	3
1.3.1	Electrical Distribution .....	4
1.3.2	Hydronic System.....	5
1.4	NREL HPC Facility Cooling Overview.....	6
1.4.1	Sandia HPC Facility Cooling Overview .....	7
1.5	Installation at NREL.....	7
<b>2</b>	<b>Test Methodology</b> .....	<b>8</b>
2.1	Data Collection.....	8
<b>3</b>	<b>Results</b> .....	<b>10</b>
3.1	Discovery Round of Testing .....	10
3.2	Long-Duration, 20-Hour Data Set.....	11
<b>4</b>	<b>Conclusion</b> .....	<b>14</b>
<b>5</b>	<b>Suggestions for Further Work</b> .....	<b>15</b>
	<b>References</b> .....	<b>16</b>

## List of Figures

Figure 1. Graphic of Aquarius fixed cold plate and server trays .....	1
Figure 2. Conceptual design of the Aquarius cooling technology .....	3
Figure 3. Rack hardware (front view, back view).....	4
Figure 4. Cooling system schematic for the HPC Data Center located in the ESIF. ....	6
Figure 5. Diagram, metering locations.....	9
Figure 6. Power data from Powershelf, supplying compute nodes, for 20-h data set.....	11
Figure 7. Facility Btu meter data for 20-h data set .....	12

## List of Tables

Table 1. Discovery Round of Testing .....	10
Table 2. Mean Energy Flow Rates for 20-hour data set .....	13




# 1 Introduction

The National Renewable Energy Laboratory (NREL) in Golden, Colorado, and Sandia National Laboratories in Albuquerque, New Mexico—two U.S. Department of Energy national laboratories—have partnered on recent projects deploying high-performance computing (HPC) liquid-cooled systems and working through energy-efficient data center designs. Both have been learning from one another as they work toward having the ultimate energy-efficient data center with the least environmental impact.

In the first half of 2018, Aquila installed its fixed cold plate, liquid-cooled Aquarius rack solution for HPC clustering at NREL’s Energy Systems Integration Facility (ESIF) as part of a partnership with Sandia. Aquila and Sandia chose NREL’s HPC Data Center for the initial installation and evaluation because the data center is configured for liquid cooling and has the required instrumentation to measure flow and temperature differences to facilitate testing. This cluster will be returned to service at Sandia’s recently completed warm-water-cooled HPC data center in early 2019.

The objective of the joint Sandia/NREL project was to evaluate the Aquarius fixed cold plate, warm-water cooling technology in a real-world liquid data center environment. Water as a heat-exchange medium is three orders of magnitude more efficient than air, and getting the heat exchange closest to where the heat is generated is most efficient. This new fixed cold plate, warm-water cooling technology (reference Figure 1) together with a manifold design also provides easy access to service nodes and eliminates the need for server auxiliary fans altogether. This paper gives an overview of the Aquarius fixed cold plate cooling technology and provides results from early energy performance evaluation testing.



**Figure 1. Graphic of Aquarius fixed cold plate and server trays**

## 1.1 Brief Aquarius Development History

In June 2015, Aquila entered into a technology transfer and licensing partnership to develop and productize Clustered Systems Company’s patented fixed cold plate refrigerant cooling system toward a warm-water-cooled HPC rack system addressing commercial, off-the-shelf, X86\_64 clustered computing. A successful prototype was powered up in June 2016, and the production

design and initial HPC system was assembled, tested, and declared ready for customers by the summer of 2017. The Clustered Systems fixed cold plate approach cools all the major heat sources. The Aquarius warm-water rack is a third-generation design built on 10 years of research-and-development that delivered two rack systems along the way. Sandia purchased the first Aquarius unit from Aquila for evaluation and testing. For more on the development history, see the [Chill-Off 2](#) (Coles 2010) final report and [ExaBlade](#) final report (Hughes and Lipp 2013).

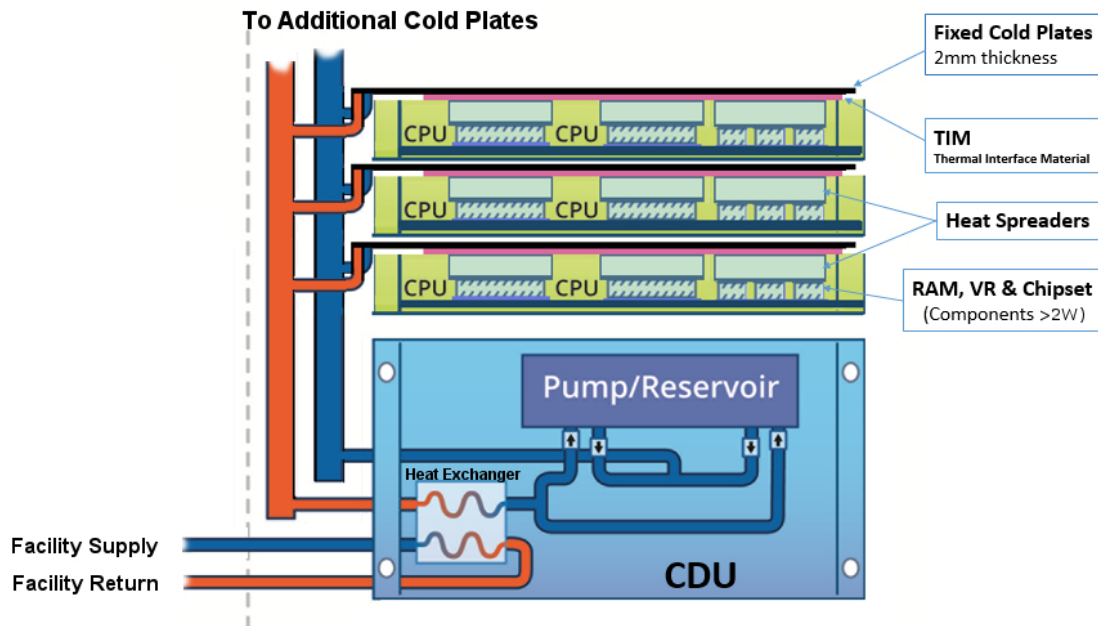
## 1.2 Aquila's Aquarius Cooling System

The Aquarius system uses several innovations to provide an energy-efficient touch cooling-derived rack system addressing commercial, off-the-shelf dual processor server systems.

Critical technology developments of the warm-water, fixed cold plate system are as follows:

- Horizontally mounted water-driven fixed cold plates and manifold system
- Compliant Thermal Interface Material (TIM) specially developed with both stiffness and compliancy, in order to eliminate planarity issues and transfer heat efficiently to the fixed cold plate
- Server tray with reliable lift mechanism to make firm contact with the fixed cold plates
- 12-VDC Hot Swap Power Board (HSPB) compensating for rise and fall times to ensure DC power bus stability during servicing
- Intelligent Management Platform Interface (IPMI) control circuitry with an Inter-Integrated Circuit (I<sup>2</sup>C) circuit on the HSPB allowing for remote control of the node motherboards as well as granular power measurement and monitoring of all IPMI reporting data points
- System-compliant compute node faceplates with lit power button and activity lights displayed. (Because there are no fans, the nodes are so quiet that visual feedback is needed to know if the system is powered up.)

The cold plates are literally a sandwich of two stainless steel plates with weld points inside to distribute turbulent flow within them. The flow pattern ensures a consistent temperature across each 21-in. by 28-in. cold plate. The cold plates are sealed with a laser welding process and pressure tested for leaks. The manifold delivers water to each plate by way of a small copper elbow between the large intake manifold and each cold plate. Metals used are “water friendly,” and no dissimilar metals are used that could cause corrosion. Each cold plate is 2-mm thick and contains about a pint of water. A 12-plate manifold was selected for Sandia's HPC cluster. (Note: Up to 24 cold plates can be accommodated in a 42-open unit [Open Compute Project](#) [OCP] rack.) Flow is regulated at 1 gpm at each of the cold plate copper intakes, and the 12-plate manifold delivers a max of 12 gpm. Flow and pressure can be adjusted slightly up or down depending on the thermal output in watts of the server. In the configuration delivered, each cold plate can remove >2 kW of heat with 1 gpm of total flow at an input temperature of 85°F.



**Figure 2. Conceptual design of the Aquarius cooling technology**

All components that generate more than 2 W have their heat removed through aluminum heat spreaders mounted on the central processing units (CPUs), random-access memory (RAM) dual in-line memory modules (DIMMs), and chipsets so that they are in touch contact with the warm-water, fixed cold plate, as shown in Figure 2. The small amount of heat generated by devices less than 2 W is shed through the copper in the motherboard and out through the metal chassis to data center environment. A compliant thermal interface material (TIM) sits between the heat spreaders and the fixed cold plate. The TIM eliminates any potential planarity issues that routinely occur because of minor component height variance, motherboard layout, or board manufacturing process inconsistencies.


### 1.3 Rack Hardware

Aquila focused on using the OCP geometry for the initial rack systems as:

- The 21-in. mounting width allowed for three 6.8-in.-wide server trays to be cooled under each fixed cold plate.
- The 1.88-in.-open unit height allowed room for insertion and extraction of 1.75-in. high trays. Each tray includes a mechanical mounting system that, when latched, pulls the tray to the cold plate, assuring good contact for heat removal.
- High server density allows for close coupling of servers to high-bandwidth, low-latency Mellanox or Intel Omni-Path edge switches using inexpensive copper cabling.

Sandia’s Aquila-based HPC cluster was named “Yacumama” and was configured to operate independently from all other HPC systems in the ESIF data center. There are 36 compute nodes that have INTEL S2600KP motherboards. The motherboards are configured with dual X86\_64 XEON CPUs, 128 GB of RAM, 128 GB solid-state drive (SSD), and an Omni-Path adapter.

The head node is configured with the same CPUs and RAM count as the compute nodes and boots from a small Universal Serial Bus drive resident inside the head node. The head node is where programs are launched to the compute nodes for parallel processing. Message passing occurs over the Omni-Path switch in a fat tree topology. All nodes are fully IPMI 2.0 compliant and can be powered up and shut down via remote commands. On paper, the configuration supplied would be capable of providing >40 teraflops of LINPACK performance while drawing less than 15 kW of power.



**Figure 3. Rack hardware (front view, back view).**

Roman numerals I, II, and III identify hardware groupings, used for reporting results (Section 3).

### 1.3.1 Electrical Distribution

This initial Aquarius rack design uses the OCP RackV1 for the rack geometry and power distribution design. The OCP 12-VDC bus rail system allowed Aquila to design a dense power delivery system, enhancing the 42-open unit rack’s power efficiency and maximizing off-the-shelf server density. The architecture matched well with the compute node’s power requirements and provided more than adequate liquid cooling. Power is delivered by a bulk redundant 415-VAC to 12-VDC power supply. Its nine rectifiers can deliver up to 21.6 kW of redundant power at 415 VAC. The use of bulk redundant AC-to-DC power conversion eliminates a step down in the information technology (IT) power chain, saving about 5% of total load. Each server pulls its power through the 12-VDC bus bar rail system; and each power zone, in this case, can provide power for up to and beyond the 36 INTEL S2600KP motherboards. This method provides for a good deal of power headroom to absorb spikes typical of certain HPC applications, such as the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS). The OCP rack system allowed Aquila to design a dense power delivery system maximizing the server density. This 42-open unit rack system cools as many as 72 S2600KP server nodes simply by adding a second liquid insert and manifold to the top half of the cabinet.


### **1.3.2 Hydronic System**

Most liquid-cooling approaches involve a cooling distribution unit (CDU), which interfaces with the facility cooling loop and provides cooling liquid at the appropriate temperature, pressure, and chemistry for the IT equipment cooling loop (secondary closed loop).

The 4U rack CDU was custom designed to meet the Aquarius warm-water flow and pressure specifications for 230-VAC, single-phase, 60-Hz power input. Operating power draw is approximately 750 W to 800 W. This pumping power easily cools the heat generated by the 36 compute nodes, transferring the heat to the facility warm-water loop at NREL. For multiple rack installations, end-of-row rack CDUs can be deployed to distribute and manage flow to the Aquarius liquid manifolds.

## 1.4 NREL HPC Facility Cooling Overview

In building its data center, NREL’s vision was to create a showcase facility that demonstrates best practices in data center sustainability and serves as an exemplar for the community. The innovation was realized by adopting a holistic “chips to bricks” approach to the data center, focusing on three critical aspects of data center sustainability: (1) efficiently cool the IT equipment using direct, component-level liquid cooling with a power usage effectiveness (PUE) design target of 1.06 or better; (2) capture and reuse the waste heat produced; and (3) minimize the water used as part of the cooling process. There is no compressor-based cooling system for NREL’s HPC data center. Cooling liquid is supplied indirectly from cooling towers. Direct component-level liquid cooling is accomplished with up to a 75°F (24°C) cooling water supply, and the return water containing the waste heat is 95°F–105°F (35°C–40°C).



**Figure 4. Cooling system schematic for the HPC Data Center located in the ESIF.**

Heat rejection options: (1) energy recovery/building heat, (2) thermosyphon, (3) cooling towers.

All heat energy from the data center IT equipment is captured in the energy recovery water (ERW) closed-loop system. The three heat-rejection options for this IT load operate according to the following hierarchy:

1. When possible, heat energy from the energy recovery loop is transferred through the energy recovery heat exchanger to the building process hot water loop to help heat the building or campus.

2. After reuse potential is exhausted, warm ERW flows to the fourth-floor mechanical space. When temperatures permit, heat is dissipated through the thermosyphon cooler (an advanced dry cooler that uses refrigerant in a passive cycle to dissipate heat).
3. The remaining heat is transferred from the ERW loop to a tower water open loop via the cooling tower heat exchanger.

NREL's HPC data center [website](#) contains more background information.

#### **1.4.1 Sandia HPC Facility Cooling Overview**

Sandia's new HPC facility was a design/build data center built in 2018 and obtained a Leadership in Energy and Environmental Design Gold certification. The design for power will expand to 14 MW, is 85% direct warm-water cooled with 75°F–85°F (24°C–29°C) entering, and is 15% air-cooled using airside economization. The facility is fully automated and features a 14,000-ft<sup>2</sup> (1,300-m<sup>2</sup>) computer room floor and novel energy- and water-saving infrastructure technologies. This includes a new thermosyphon technology previously validated for use with HPC systems through a [joint NREL/Sandia research project](#). The thermosyphon uses a passive refrigeration cycle that provides inexpensive liquid cooling without the excessive water use required with more common evaporative cooling approaches. When fully occupied, the facility will save up to ~15 million gallons (57,000 m<sup>3</sup>) of water per year when compared with evaporative cooling use. This facility enables Sandia to remain at the forefront of computational engineering and science and a leader in HPC research.

### **1.5 Installation at NREL**

The Yacumama system was shipped to NREL on three pallets. The OCP rack with CDU, manifold, cold plate system, and Powershell already installed was rolled off the first pallet. Because this rack was shipped in winter, and with possible subzero temperatures in Colorado, it was filled with ethylene glycol to avoid any freezing potential inside the rack's cooling manifolds and cold plate system. The glycol was pumped out on-site, and the secondary IT closed loop was then filled with an Aquila-supplied mixture of soft water, biocide, and a corrosion inhibitor. This formulation was recommended by the CDU supplier to provide the highest cooling efficiency. Although using glycol in the secondary IT coolant loop is acceptable, it is not required in the design, and using glycol decreases the heat transfer efficiency in a small yet measurable way. The other two pallets contained the 36 compute nodes and head node, which were unpacked and slid into the rack.

The facility warm-water loop was brought from the ESIF water distribution manifold to the Aquarius rack via 1-in. insulated rubber hoses. These were attached to the CDU's primary loop inlet and outlets. Line power was provided from a row power distribution unit (PDU), with one feed for the Powershell and another for a rack PDU used to power the CDU, switches, and head node.

The installation was straightforward and completed in half a day. The only required procedure on the IT manifold loop was to properly flush the manifold, then refill and bleed the manifold/plates.

## 2 Test Methodology

The main objective of the testing conducted was to determine the quantity of energy being dissipated to water to determine how close the fixed cold plate approach is to a room-neutral solution. To accomplish this, measurements were taken under constant max load conditions using the LINPACK benchmark. An initial day of discovery testing was first conducted to ensure the fixed cold plate design performed to the specification and to get familiar with the system at NREL's HPC center environment. This was a phased approach with measured data/metrics to establish a baseline before changing the conditions to the system. In theory, the system should perform at the same level throughout all changed parameters with no recorded failures because of thermal limitations. Note that facility conditions could not be altered for any testing, and the facility supply water was typically around 72°F (22.2°C). The best attempt was made during testing to reach a semi-steady state for estimating energy recovered using facility-caliber instrumentation.

- Discovery testing:
  - The purpose was to identify the range of reasonable changes that can be made on the system—from supply temperature to fixed cold plates and the impact of adjusting the facility flow. We initially recorded the idle state readings and then ran LINPACK tests at 100% CPU use while varying the supply temperature to the cold plates and adjusting the CDU settings. Finally, we performed a ramp-up test (idle to full power) to find the highest supply temperature that we could identify as safe for the system. The goal was to achieve a 10-minute window under each test case where the energy rate measured at NREL's ERW return manifold was relatively consistent under a constant load.
- Long-duration test:
  - After discovery testing led to a comfortable IT supply temperature, then LINPACK tests at 100% duty cycle were run for 48 hours. The data logging system needed to be restarted midway, so this paper reports on a 20-h window in which the system was held at a semi-steady state for reporting time-averaged data. Note that for the LINPACK 100% duty cycle runs, memory allocation was set at 82.5%.

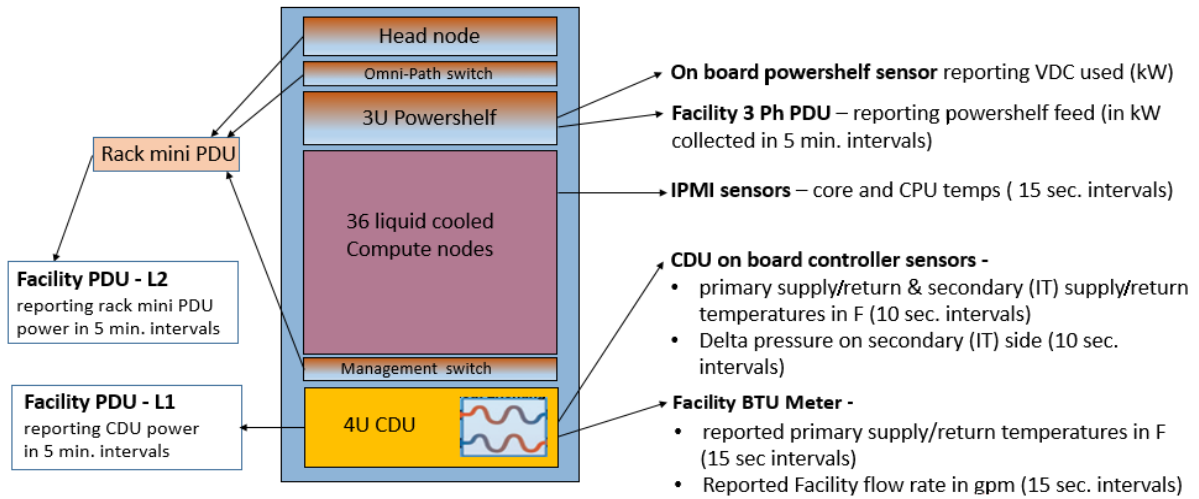
### 2.1 Data Collection

Key data points:

- Monitoring power (415 VAC) delivered to Powershelf
- Monitoring power from Powershelf sent to compute nodes (accounts for power conversion loss of 415 VAC to 12 VDC)
- Monitoring power (230 VAC) to other equipment in rack (CDU, head node, and two switches)
- Facility manifold, recording energy flow rate (heat) captured direct to facility ERW, assume remaining energy flow rate transferred to air
- Rack CDU recording IT loop inlet/outlet temperatures, facility inlet/outlet temperatures



- IPMI 2.0 sensor data reporting core temperatures of CPU 0 and 1, along with system health/status.



**Figure 5. Diagram, metering locations**

## 3 Results

This section provides the test results for the discovery phase and 20-h data set.

### 3.1 Discovery Round of Testing

General discovery round test notes:

- Test 1: The Yacumama system was idle prior to the test start. The mean facility flow is not reported in Table 1 because there were wide swings in flow rates as a result of the CDU facility control valve showed hunting behavior (before CDU adjustments were made).
- Test 2: At 14:00, the LINPACK 100% was launched and continued through Test 5.
- Test 3: The CDU set point was increased, and we adjusted the CDU deadband prior to the 10-minute window.
- Test 4: The CDU set point was increased, and we adjusted the CDU deadband prior to the 10-minute window.
- Test 5: The CDU set point was increased prior to the 10-minute window.
- Test 6 (next day): The Yacumama system was idle prior to the start, and the final CDU integral time and deadband settings were entered prior to the test start. At 10:22, the LINPACK 100% was launched and continued through the end of the 10-minute window.
- All tests: The mean facility supply temperature for all discovery round tests was between 72.2°F–72.8°F (22.3°C– 22.7°C).

**Table 1. Discovery Round of Testing**

Test Number Timeframe	Test Condition	CDU IT Supply Set Point, °F (°C)	Mean Power to Compute Nodes (kW)	Mean Facility Flow (gpm)	Mean Temp. Facility Return °F (°C)	Mean Temp CPUs °F (°C)	Max Temp CPUs °F (°C)
Test 1 13:45–13:55	Idle state	80 (26.7)	3.1	See note	77.5 (25.3)	87.3 (30.7)	95 (35)
Test 2 14:10–14:20	100% LINPACK	80 (26.7)	12.6	10.6	80.4 (26.9)	145.4 (63.0)	165 (74)
Test 3 14:40–14:50	100% LINPACK	85 (29.4)	12.7	6.2	87.2 (30.7)	151.7 (66.5)	172 (78)
Test 4 16:00–16:10	100% LINPACK	90 (32.2)	12.7	4.5	92.5 (33.6)	156.9 (69.4)	178 (81)
Test 5 16:30–16:40	100% LINPACK	95 (35.0)	12.6	3.4	98.5 (36.9)	162.7 (72.6)	183 (84)
Test 6 11:10–11:20 (next day)	Ramp up to 100% LINPACK	95 (35.0)	12.6	3.4	98.2 (36.8)	162.3 (72.4)	183 (84)

## 3.2 Long-Duration, 20-Hour Data Set

General notes for the 20-h data set:

- Time span: 18:00–14:00 (*following day*)
- Test condition: 100% LINPACK
- CDU IT supply set point: 85°F (29.4°C)
  
- Energy recovered direct to water: 251.1 kWh
- Mean facility return temperature: 87.1°F (30.6°C)
- Mean facility supply temperature: 71.4°F (21.9°C)
- Note: The facility delta temperature from the CDU was used because those temperature sensors are located closest to the rack (whereas the Btu meter temperature sensors are located farther away).
  
- Mean of all CPU-1 temperatures: 160.7°F (71.5°C)
  - Min. individual node mean: 148.9°F (64.9°C)
  - Max. individual node mean: 179.4°F (81.9°C)
- Mean of all CPU-2 temperatures: 152.4°F (66.9°C)
  - Min. individual node mean: 139.8°F (59.9°C)
  - Max. individual node mean: 161.2°F (71.8°C)
- Also of note: Turbo was enabled full time in the firmware, and no CPU clock throttling was ever observed (including during the discovery round of testing). This was fully expected by the Aquila team based on their previous lab testing and was verified at the ESIF.


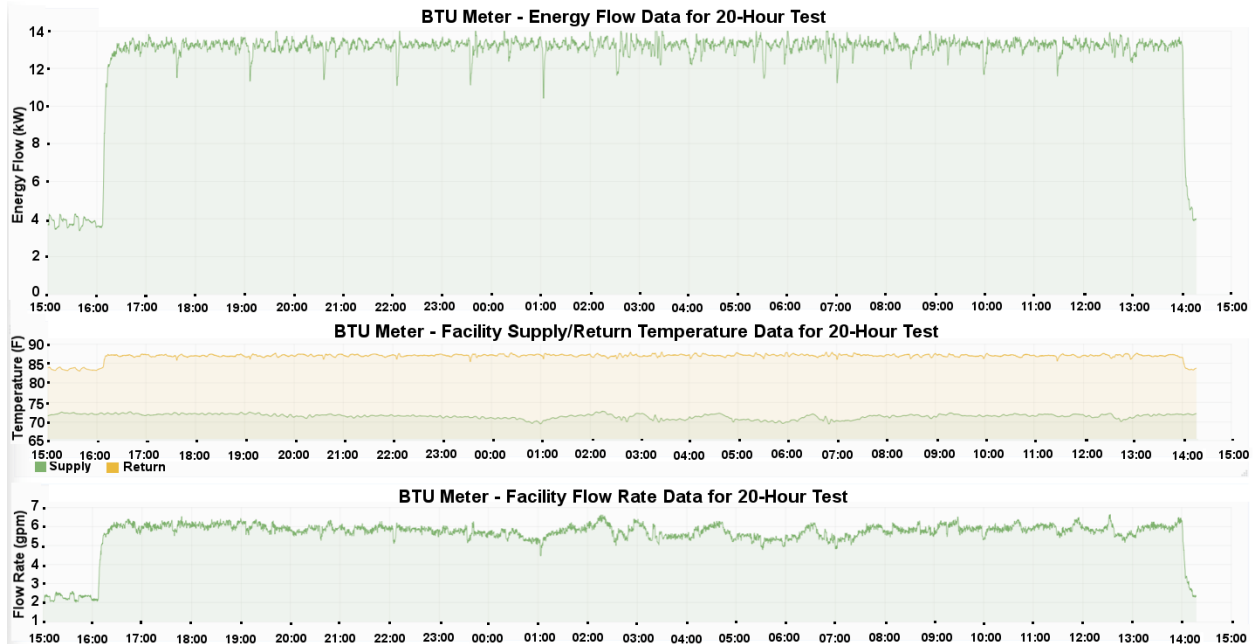


Figure 6. Power data from Powershell, supplying compute nodes, for 20-h data set



**Figure 7. Facility Btu meter data for 20-h data set**

The percentage of heat captured direct to water as reported in Table 2 is based on hardware groupings as shown in Figure 3 and stated as:

1. Compute nodes only
2. Compute nodes and Powershell
3. Compute nodes, Powershell, CDU, switches, and head node.

A high percentage of heat energy recovered direct to water is desired for a system to be room neutral. That saves on facility capital costs along with operating costs when a rack-based cooling system is scaled up in a data center. Any heat that is not captured at the rack level must be handled by the facility system (in the ESIF's case, by the use of fan walls, which are not as efficient).

**Table 2. Mean Energy Flow Rates for 20-hour data set**

<b>By Hardware Grouping</b>	<b>Heat Transfer Rate to Water (kW)</b>	<b>Power for Hardware Grouping (kW)</b>	<b>Energy Recovery Direct to Water (Percentage)</b>
I: Compute nodes only	12.6	12.8	98.3% <sup>a</sup>
II: Compute nodes and Powershell	12.6	13.4 <sup>b</sup>	93.4% <sup>a</sup>
III: Compute nodes, Powershell, CDU, switches, head node	12.6	14.5	86.1%

<sup>a</sup> Does not take into account the portion of CDU energy that is transferred to facility water—there was not sufficient instrumentation in place to determine this. This percentage represents the upper limit.

<sup>b</sup> Recorded data showed Powershell efficiency was 95%, so there was 0.6 kW of heat energy being transferred to air.

## 4 Conclusion

We found that a very high percentage of heat was captured direct to water. And by eliminating the use of auxiliary fans, all the compute node's power draw is used for computing (no stranded fan energy). Ideally, the fixed cold plate solution could be expanded to support the Powershell, head node and switches in future system designs.

The system was in operation at the ESIF data center for nearly 10 months, requiring zero maintenance, and no water leaks were observed. The manifold and cold plate design are observably robust in nature.

Based on observations, it seems possible to increase the supply IT water temperature to the compute nodes. This would allow warmer facility supply as well as push the return temperatures up, which is ideal for waste heat reuse and/or use of dry coolers for on-site water savings.

## 5 Suggestions for Further Work

We suggest the following if further testing is conducted:

- Move the Btu meter (flow meter with temperature sensors inserted into the facility's supply and return lines) as close to the rack CDU as possible
- Add thermocouples to the secondary loop side to verify the supply temperature to the fixed cold plates from the CDU.
- Perform a flow meter check (gravimetric calibration, timed flow into a tank), or check against a second calibrated flow meter (in addition to having matched temperature sensors for Btu meter).
- Use a power quality meter to measure the feed to the Powershell, allowing for a higher sample rate than a PDU typically provides.
- Review the energy losses in the 12-VDC bus rail system based on thermal images taken during the discovery round of testing. Power reaching the compute nodes is likely a bit lower than what is reported by the Powershell.

## References

Coles, H.C. 2010. *Demonstration of Alternative Cooling for Rack-Mounted Computer Equipment*. Berkeley, CA: Lawrence Berkeley National Laboratory.

[https://datacenters.lbl.gov/sites/default/files/Demo%20Alternative%20Cooling\\_Rack%20Mounted\\_2010.pdf](https://datacenters.lbl.gov/sites/default/files/Demo%20Alternative%20Cooling_Rack%20Mounted_2010.pdf).

Hughes, Philip N., and Robert J. Lipp. 2013. *Recovery Act: Development of a Very Dense Liquid Cooled Compute Platform*. Washington, D.C.: U.S. Department of Energy, Office of Scientific and Technical Information. <https://www.osti.gov/servlets/purl/1110242>.