Operational Data Analytics: Optimizing the National Energy Research Scientific Computing Center Cooling Systems

Norman Bourassa
njbourassa@lbl.gov
Lawrence Berkeley National Laboratory

Walker Johnson, P.E., LEED AP
wjohnson@kw-engineering.com
kW Engineering, Inc.

Jeff Broughton
jbroughton@lbl.gov
Lawrence Berkeley National Laboratory

Deirdre McShane Carter, PE
dmcarter@lbl.gov
Lawrence Berkeley National Laboratory

Sadie Joy, P.E.
sjoy@kw-engineering.com
kW Engineering, Inc.

Raphael Vitti, P.E., BEAP
rvitti@kw-engineering.com
kW Engineering, Inc.

Peter Seto
jordanruthseto@aol.com
Energy Efficient HPC Working Group

ABSTRACT

In 2017/2018, the Energy Efficient HPC Working Group (EE HPC WG) Dashboard Team conducted an analysis that assessed the current use of information dashboards for operational facility management in major supercomputing centers around the globe, resulting in the formalization of a process now referred to as Operational Data Analytics (ODA). Subsequent to surveys of multiple HPC facilities, the EE HPC WG determined that case studies were needed to help HPC facilities evaluate the value of implementing ODA practices. This paper provides a summary of the successful use of an ODA approach being used by the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (Berkeley Lab) in meeting organizational energy efficiency performance goals for the NERSC HPC cooling systems.

CCS CONCEPTS

• **Applied computing → Enterprise data management**; **Data centers**; • **Mathematics of computing → Time series analysis**; *Exploratory data analysis*; • **Hardware → Power and energy**; *Enterprise level and data centers power issues*; • **Information systems → Business intelligence**; *Data analytics.*

KEYWORDS

building management systems, building controls, energy efficiency, data centers, operations, high-performance computing, data collection, data analytics, time series data, Green HPC, PUE, iTUE

ACM Reference Format:

Norman Bourassa, Walker Johnson, P.E., LEED AP, Jeff Broughton, Deirdre McShane Carter, PE, Sadie Joy, P.E., Raphael Vitti, P.E., BEAP, and Peter Seto. 2019. Operational Data Analytics: Optimizing the National Energy Research Scientific Computing Center Cooling Systems. In *48th International Conference on Parallel Processing: Workshops (ICPP 2019), August 5–8, 2019, Kyoto, Japan.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3339186.3339210

ICPP 2019, August 5–8, 2019, Kyoto, Japan

ACM ISBN 978-1-4503-7196-4/19/08…$15.00
https://doi.org/10.1145/3339186.3339210

1 INTRODUCTION

This case study documents an effective use of operational data instrumentation, analysis, integration, and archiving, toward effective design, commissioning, and optimization of power usage effectiveness (PUE [1]) in high performance computing (HPC) facility environments. The Energy Efficiency HPC Working Group (EE HPC WG) refers to this process as Operational Data Analytics (ODA). In 2017/2018, the EE HPC WG Dashboard Team conducted an analysis that assessed the current use of information dashboards for operational facility management in major supercomputing centers around the globe. Most of the surveyed sites provided specific answers about their facility related systems. A few early adopter sites answered that they either had aggregated or were attempting to aggregate both facility and HPC system data. The resulting survey work led to the creation of an EE HPC WG ODA Team whose ongoing objective is to identify how these early adopters are using their systems with specific use-case scenarios. Additional investigations looked into implementation issues, such as system-to-system interfacing; the question of known or potential scalability constraints will also be documented by the ODA Team.

A particularly sophisticated ODA approach is used at the National Energy Research Scientific Computing Center (NERSC)[1] at Lawrence Berkeley National Laboratory (Berkeley Lab)[2]. The companion paper by E. Bautista et al [3] describes the NERSC-developed Operations Monitoring and Notification Infrastructure (OMNI) system, which powers their ODA process. The NERSC Energy Efficiency Team (EE team) has been utilizing ODA tools to continuously improve facility operations since their initial 2015 occupancy of a

[1] About NERSC. October 10, 2018. https://www.nersc.gov/about/
[2] About the Lab. Retrieved May 13, 2019 from https://www.lbl.gov/about/

new Berkeley Lab-located HPC facility named Shyh Wang Hall.[3] These ODA tools and processes have been critical for maximizing reliability and meeting organizational performance goals for the EE team, on which the Berkeley Lab and kW Engineering authors are active participants.

ODA has been critical at NERSC since the cooling plant systems do not use compressor-based cooling. Instead, NERSC employs direct-evaporative cooling, air-side economizer modes, and cooling tower generated cooling water (CW). The first two can greatly affect the supply air humidity and cannot always meet the preferred supply air temperature setpoints. As a result, facility operations staff rely on precise monitoring of temperature, humidity, and HPC operational data of the computer room to inform system control settings, record performance data, and learn improved operations toward maximized reliability.

Another important outcome of this sophisticated ODA approach is ongoing EE optimization, which supports Berkeley Lab's federal, state, and University of California energy efficiency and PUE goals and requirements. With a team available for ongoing monitoring of the Shyh Wang Hall energy performance and the available ODA systems, EE opportunities can be identified and implemented with a high level of confidence and operational safety for the HPC equipment.

Since most data center facilities typically have an HPC machine refresh rate of 3 to 5 years, commissioning new HPC machines, facility cooling infrastructure, and ongoing operations are merging into one continuous process. At NERSC, the ODA functions have become an integral part of the Shyh Wang Hall ongoing operations.

2 NERSC FACILITY OVERVIEW

Jeff Broughton, Deputy for Operations at NERSC, and his team have logged two decades of practical experience operating energy-efficient HPC facilities, first at NERSC's downtown Oakland Scientific Facility (OSF) and now in Shyh Wang Hall. The main HPC systems in use at NERSC are the hybrid liquid and air-cooled Cori[4] and Edison[5] Cray XC Series®supercomputers, multiple air-cooled HPC clusters, and a high-performance storage system (HPSS).[6] The Cray XC cabinets are configured in cooling air stream linked rows, with dedicated blower fan cabinets powering the air flow. Within each Cray XC compute cabinet are "backdoor style" cooling coils that extract waste heat to the facility cooling water (CW) loop. The CW loop is connected to the closed loop side of liquid-to-liquid heat exchangers which connect with an open loop tower water (TW) pumping system that rejects the heat to the outside air with cooling towers. The entire computer room air and air-cooled systems are cooled by air handling units that utilize air-side economizers, direct evaporative coolers, and CW cooling coils.

[3]Weiner, Jon. Berkeley Lab Opens State of the Art Facility for Computational Science. November 12, 2015. https://newscenter.lbl.gov/2015/11/12/facility-for-computational-science/

[4]Cori, NERSC Systems. May 2019. https://www.nersc.gov/systems/cori/

[5]Edison, NERSC Systems. May 2019. https://www.nersc.gov/systems/edison-cray-xc30/ This system was decommissioned on May 13, 2019.

[6]The two NERSC HPSS systems are the sole exceptions to compression-free air conditioning, since the tape archive drive cabinets are decoupled from the general computer room air and include integral DX cooling and humidity control units. https://docs.nersc.gov/filesystems/archive/

3 ONGOING MONITORING AND OPTIMIZATION TOOLS

NERSC leverages multiple monitoring systems for ongoing monitoring and optimization, with the two most prominent being OMNI and SkySpark.[7] As described in Bautista et al [3], OMNI is not a single tool, but a versatile platform of applications that combine a vast amount of HPC and IT systems data with comprehensive cooling and facility systems performance data. SkySpark is ideal for interfacing with building systems (e.g. via BACnet), collecting and analyzing the available building data, and setting up virtual, calculated points derived from that data. By eliminating the work associated with downloading and processing data in more typical engineering offline tools (such as spreadsheets), the process of ongoing optimization of NERSC operations can be performed much more efficiently.

The NERSC OMNI system has merged the Building Management System (BMS), an array of supplemental rack-level IT sensors, and Cray syslog data into a real-time, searchable, and easily graphed ODA system. The resulting Elasticsearch database absorbs and archives 25k data points/sec at present, with plans to expand it to as many as 100k data points/sec in late 2019. All this operational performance data has been archived within the general HPSS, dating back to the initial Shyh Wang Hall commissioning and occupancy. Rapid graphing of the real-time data is principally provided by the Elasticsearch Grafana[8] and Kibana[9] web browser-based user interfaces, while numerous other analysis and data visualization needs are filled with more specific open-source software tools when needed. [2]

The SkySpark platform, used campus-wide at the Berkeley Lab, gathers data through a live link to the NERSC BMS, the Elasticsearch database, and an ION power meter database. SkySpark provides an environment for custom analysis. The EE team has built custom views for scatter plots, trend graphs, and performance metrics which are continuously updating with the live OMNI data link. These views allow the EE team to continuously monitor system efficiency, spot energy efficiency opportunities, and verify energy savings. SkySpark is hosted on the Berkeley Lab private network for security reasons, but it can be accessed remotely through VPN for use by contracted engineering consultants or key NERSC staff.

3.1 ODA Related Projects

The ODA process at NERSC has resulted in steady, ongoing discovery and implementation of EE and reliability measures. The EE team generated estimated energy savings upon discovery of each measure and verified the savings after implementation was complete. Table 1 summarizes the EE measures identified since 2017.

Table 2 provides a summary of the OMNI measured average PUE of 1.07 in Shyh Wang Hall from May 2018 to April 2019, which is well below the federal goal of a PUE between 1.2 to 1.4 for new facilities. Early 2019 results show indications of the recent EE project benefits driving a PUE reduction. During 2018, NERSC deployed a

[7]SkySpark, SkyFoundry, 2019. https://skyfoundry.com/product

[8]Grafana Labs. 2019. https://grafana.com

[9]Elasticsearch. 2019. https://www.elastic.co/products/kibana

Table 1: Energy Efficiency Measures Identified at NERSC

MEASURES OVERVIEW NERSC

	Measure Title	Energy Savings (kWh)		Water Savings	Cost Savings	Est. PUE
		Estimated	Verified	m^3	$	Reduction
1	Install Firmware to Enable ESS Mode for UPSs		350,000	500	$20,300	0.007
2	Implement Tower Water Supply Temperature Reset and Reduced Tower Water Pump Speed		420,000	600	$24,360	0.009
3	Reset Cooling Water Temperature Setpoint and Enable Cray Dynamic Fan Control		400,000	600	$23,200	0.008
4	Install New Heat Exchanger		780,000	1,200	$45,240	0.016
5	Install Bypass Valves		25,000	0	$1,450	0.001
6	Reset Cray Air Temperature Setpoint	200,000		300	$11,600	0.000
7	Optimize Dynamic Fan Control	200,000		300	$11,600	0.000
8	Install Booster Pump	240,000		400	$13,920	0.005
9	Install Cold Aisle Temperature Sensors and Optimize AHU SAT and Flow Control	300,000		500	$17,400	0.006
10	Install Cray Supply Air Hoods	100,000		200	$5,800	0.002
	Total	1,040,000	1,975,000	4,600	$174,870	0.054

Cost of electricity $0.058 /kwh Annual Average PUE = 1.07

rigorous 15-minute interval PUE calculation method which evaluates up to 50 metering points. The method was an assessment outcome which identified all electrical loads directly supporting the HPC infrastructure and operations, such as lighting, freight elevators, control room air conditioning, etc. One significant caveat is that the Cray XC Series systems currently have blower fan power co-mingled with the system compute power, so PUE values and the calculation method will be revised after completion of an iTUE [4] project to dis-aggregate these data points (Briefly discussed in the *Ongoing and Future Projects* section).

Table 2: Summary of Average Measured PUE at NERSC

NERSC - Monthly Average Measured PUE	
May-2018	1.07
Jun-2018	1.08
Jul-2018	1.08
Aug-2018	1.08
Sep-2018	1.07
Oct-2018	1.08
Nov-2018	1.07
Dec-2018	1.07
Jan-2019	1.06
Feb-2019	1.06
Mar-2019	1.06
Apr-2019	1.06
12-Month Average	1.07

4 POST-OCCUPANCY FACILITY COMMISSIONING

Shyh Wang Hall systems were initially commissioned before any significant HPC heat loads were present. Without the design heat load, the commissioning team couldn't demonstrate the systems operating as designed. Over time, as HPC equipment came online and HVAC systems ramped up in response, the EE team used ODA tools to monitor operations to ensure expected performance per design.

During this ODA process, the team noticed that the cooling plant pumps (cooling water and tower water) were operating at a much higher head pressure than expected, resulting in very high pumping power. Measurements indicated that the single heat exchanger in the cooling water loop was causing a very high pressure drop, likely as a result of a clog. The team decided to move forward with installing a second heat exchanger in parallel that would reduce the pressure drop, increase capacity for future expansion, add redundancy, and allow the team to clean the other heat exchanger without shutting down the cooling water system. The team also improved valve control to more evenly distribute flow and reduce maximum pumping energy. After these upgrades, power data showed a large reduction in annual pumping energy (Table 1, Item 4).

The scatter plot in Figure 1 shows total CW pumping power relative to CW flow rate. The baseline data (green) shows how pumping power increased to max power until they were pumping roughly 1,200 m^3/hour (5,000 gpm) at full speed. The post-installation data (blue) shows a very different pump curve. Not only are the pumps consuming less power for any given flow rate, but they are now able to reach much higher flow rates, helping to meet the cooling demand more effectively.

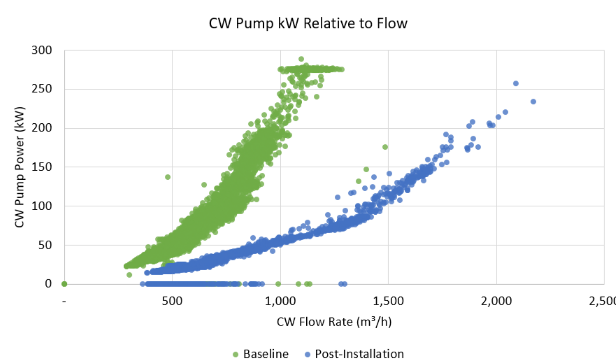

Figure 1: Pump Power Curves Before and After Installation of Second Heat Exchanger

The team also learned that the same CW pumps had to maintain a much higher pressure setpoint in order to serve a few rooftop air handler units, which provide general office space air conditioning. While the rooftop units only require a relatively small amount of actual CW flow, all the water serving the data center had to be pumped to a higher pressure to meet their CW demand. The ODA process revealed this as being a "tail wagging the dog" situation, so NERSC installed a separate and much smaller booster pump to pressurize the dedicated line serving the rooftop units. This allows the much larger main cooling plant pumps to operate at a lower head pressure.

5 ONGOING COMMISSIONING AND OPTIMIZATION

As part of Berkeley Lab campus-wide EE goals, an aggressive ongoing commissioning process has been undertaken for all campus facilities. The ODA process and tools have become indispensable for NERSC in meeting these goals, since Shyh Wang Hall is responsible for approximately 40 percent of the campus power load. The following two ongoing projects are illustrative examples of the ODA value.

6 OPTIMIZATION USING POWER-MONITORING DATA

Most of the cooling system setpoints were determined before the presence of HPC equipment loads. Once the HPC load increased to full build-out, the EE team began using ODA tools to review archived power meter data to monitor the facility cooling-plant energy performance. The primary energy-using equipment within the facility cooling plant are the cooling water (CW) pumps, tower water (TW) pumps, and cooling tower fans. The team plotted archived power data of these three equipment loads against outside air wet bulb temperature, the most important factor for the cooling tower system efficiency (Figure 2). This analysis showed that lower CW temperatures increased cooling effectiveness and allowed for reduced pumping power. The CW pumps have the highest power draw, and therefore energy consumption, during most outside air conditions. However, their power draw dropped dramatically with lower outside air temperatures. At these lower outside air temperatures, the tower fans and TW pumps operate at minimum control speeds, causing the CW temperature to drop below minimum setpoint. This reduction in temperature increased the cooling effectiveness of the Cray cooling coils, allowing the CW valves to reducing CW flow.

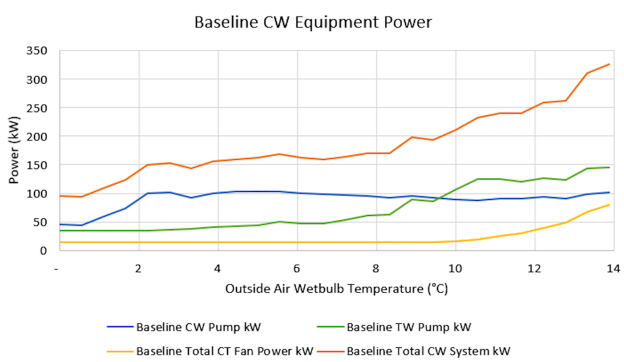

Figure 2: Major Cooling Plant Equipment Avg. Power relative to Outside Air Wet Bulb Temperature

This analysis presented a clear argument that reducing the CW temperature would reduce overall CW flow and thereby save CW pump energy, but it was unclear exactly how much cooler it should be to optimize flow for energy efficiency without overcompensating. Access to these ODA tools in real time greatly facilitated the experimental testing of this power demand hypothesis. The development of a calibrated energy model - a more traditional energy engineering method - was not necessary since the EE team could adjust various TW temperature setpoint control settings and watch the impact on various power-demand points through the ODA tools.

As a first step towards reducing the CW temperature, without over-burdening the cooling towers, the team tested a TW supply temperature reset based on outside air wet bulb temperature. This common control strategy adjusts the TW setpoint automatically as the wet bulb temperature changes. Since cooling towers reject heat primarily by evaporating water, they can never reduce the water temperature below the outside air wet bulb temperature. By adjusting the TW setpoint so that there is always a constant temperature differential (also known as "approach") between the wet bulb temperature and the setpoint, the cooling towers are able to reduce water temperature during cool weather, without resulting in excessively high fan power during warmer weather periods. The EE team determined safe minimum and maximum supply temperatures, and within that established range the TW temperature setpoint is kept at a constant approach. Figure 3 shows the total CW plant power for the measured baseline with a TW setpoint at a constant 20°C (68°F), then two experimental settings using a 2.2°C (4°F), and 2.8°C (5°F) approach versus the outside air wet bulb temperature. The graph shows that using a 2.8°C setting (green line) resulted in the lowest power draw, so it was chosen. This process is repeated as often as needed as systems are upgraded and controls change, another core benefit of the NERSC ODA tools and process.

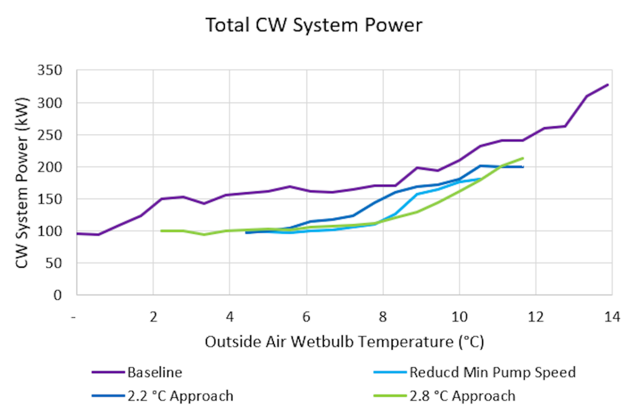

Figure 3: Total CW Plant Power setting options relative to Outside Air Wet Bulb Temperature

7 STREAMLINING THE OPTIMIZATION PROCESS USING SKYSPARK

The above process of controls optimization by means of monitoring incremental setting changes, was relatively efficient, but must be performed on an on-going basis in order to maintain high efficiency energy performance as facility churn occurs. Additionally, it often requires manual data downloads and data processing, which can be slow and labor-intensive. To improve upon the process, the EE team tied the NERSC OMNI system into LBNL's campus-wide SkySpark analytics platform. This platform is ideal for creating continuously-updated graphical representations of HVAC measurement data, thereby helping the team to monitor systems performance, quickly spot problems, identify opportunities, and verify energy savings.

Figure 4 is an example of a SkySpark graphic the EE team created to replace manually created scatter plots, which are commonly used in engineering analysis. This tool allows the user to choose a time frame and date for which a milestone settings change event was initiated to determine the impact on power draw. The graph shows how the cooling water plant total power reduced significantly at lower outside air temperatures as a result of increasing the minimum tower water supply temperature setpoint, but remained similar at higher temperatures. This change allowed the cooling tower fans to ramp down at low temperatures without making an impact on pump energy. This experiment proved that increasing the minimum setpoint was a step in the right direction. The EE team then continued to increase the setpoint, with iterative result reviews, until no further improvements were detected.

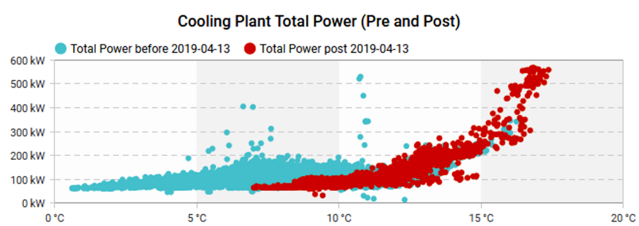

Figure 4: Two Different Total CW Plant Power results versus Outside Air Wet Bulb Temperature

The EE team is now adding HPC and IT support equipment power data to the LBNL campus SkySpark tool. This will allow the team to move beyond HVAC energy consumption and optimize total energy consumption, for example by accounting for fan power within rack mounted HPC servers or hyperscale compute cabinets. SkySpark will automatically calculate and monitor a variety of metrics, including PUE, total power utilization effectiveness (TUE), and IT power utilization effectiveness (iTUE). [4]

8 LEARNING FROM EXTREME ENVIRONMENTAL CONDITIONS

The above examples have discussed cooling water plant projects, but NERSC also has a considerable amount of HPC air-cooled systems served by air handling units (AHUs). Fortunately, due to the Berkeley climatic conditions, compressor-free cooling (air-side economizer) mode operation is possible most hours of the year, with indirect or direct evaporative cooling providing supplemental cooling during warmer periods. However, there have been occasional extreme atmospheric conditions that have stressed the facility, and ODA has helped immensely in teaching Operations how to improve the air-side control settings. An example occurred when the 2018 wildfires in Northern California[10] produced an extreme spike of particulates in the atmosphere.

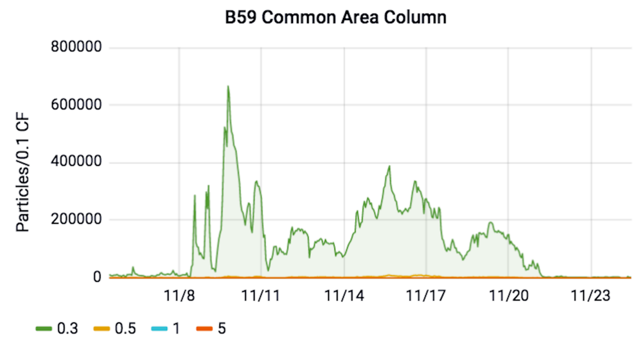

Figure 5: Computer Room Particle Counter Readings

On November 8, 2018, NERSC OMNI air particle sensors detected a rapid degradation of outdoor and indoor air quality (Figure 5), forcing operators to manually override the automated outside air damper operation into the closed position for all computer room AHUs. While these AHUs typically operate in air-side economizer mode for most of the year, when in return-air-only mode supplemental cooling was manually limited to cooling coils to avoid humidity accumulation from the direct evaporative coolers. Initially the operation looked stable, but the poor outdoor air quality continued for two weeks, and eventually multiple control-system issues began to appear. In the end, seven distinct BMS control logic fixes were identified within at least two general categories:

- Supply and mixed air damper control (temperature)
- Indoor relative humidity and condensation control

In addition, this event reinforced the need for the implementation of a "single-click" BMS feature so facility operators would be better prepared for responding to the next poor air-quality event, which will inevitably occur during the increasingly fierce annual fire seasons in California.

Almost immediately, supply air temperature control proved to be difficult since the normal BMS logic method was to control the outside air and return air dampers to a mixed air temperature setpoint. However, instability resulted with the locked outside air damper, and the control logic had to be adjusted to use the supply air temperature instead.

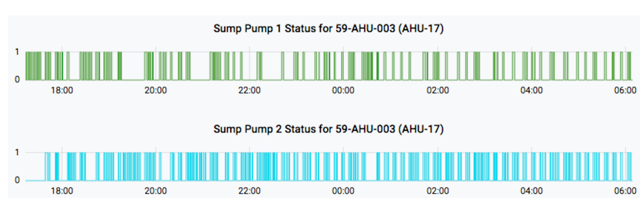

Figure 6: AHU Evaporative Cooler Pump Cycling Problems

[10]Information about the Camp Fire, Butte County, California available at Cal Fire Incidents website. http://cdfdata.fire.ca.gov/incidents/

Once the supply air temperature control problems were corrected, problems with humidity control eventually surfaced. Rapid temperature and humidity oscillations began as the evaporative cooler sump pumps cycled on and off (Figure 6) at varying rates in each AHU. Causes were traced to more conflicts in the supply air temperature setpoint logic and the direct evaporative cooler humidity (stage 1) and cooling (stage 2) staging controls. Some small and hard to detect hardware failures were also discovered. After considerable iterations between control-logic adjustments and the ODA, stable control corrections were obtained on the final day of the poor outdoor air event (Figure 7).

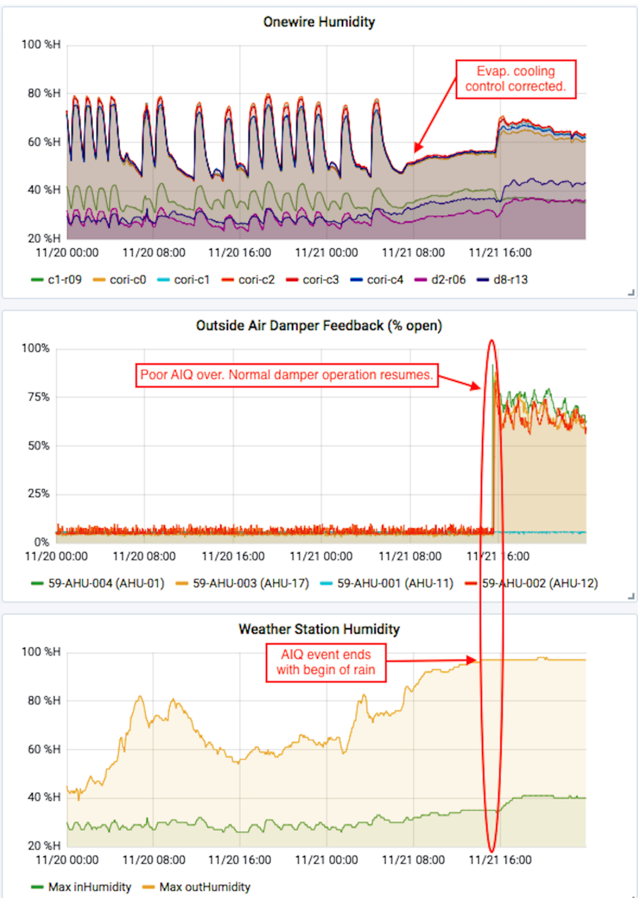

Figure 7: Supply Air Relative Humidity Control Problems and Correction

Finally, in the middle of these temperature and humidity control logic corrections, on November 16, 2018 a water-leak alarm was triggered in NERSC's computer room, and physical inspections quickly revealed that condensation was dripping off of a large CW manifold in the underfloor plenum. Since NERSC uses compression-free cooling, our general expectation had always been that condensation could not happen because the facility is tightly coupled to outside wet bulb temperatures. However, this event and the ODA process showed us that during extended periods of 100 percent return air operation, the underfloor plenum dew point temperature can indeed get above the CW temperatures produced by the cooling towers. Subsequently, operators adjusted the CW setpoint higher, and permanent logic corrections were implemented to protect against this condition in the future.

9 ONGOING AND FUTURE PROJECTS

NERSC continues to leverage its developed ODA tools to identify EE improvements as part of regular operations, and this approach is an ongoing key component of future project plans.

A project currently underway involves experiments designed to customize the performance of the Cray XC Dynamic Fan Speed Control (DFSC) feature. This feature automatically varies the cabinet blower fan speeds, based on processor temperatures, yielding reliable energy savings. It also provides systems administrators with the ability to adjust the XC Series cabinet CW cooling-coil air temperature setpoint settings. Using ODA tools, NERSC is currently developing a cooling plant communications link with the DFSC feature, which will provide the needed feedback to shift the Cray cooling systems load ratio between blower fans or the CW loop, depending on the outdoor environmental conditions imposed on the facility cooling plant. In the near-term, the team will use SkySpark to monitor iTUE and other metrics to towards a better management balance between Cray fan power and cooling plant power.

A longer term direction is to apply machine learning to optimization of plant operation. The ODA system will provide years of data that can be used as training data sets. Machine learning will help identify the system functions associated with key operating parameters, and may ultimately be applied to direct plant operations in real time.

10 CONCLUSIONS

With ongoing changes to IT and HPC equipment, high power densities, and unique cooling system configurations, ODA capabilities have helped NERSC ensure both reliability and persisting EE benefits. NERSC has demonstrated how ODA can be an essential tool for compressor-free cooling HPC facilities, and that this ODA value proposition scales very well with both machine size and the degree of HPC and facility data instrumentation integration. NERSC has also demonstrated that ODA use in experimenting and analyzing both HPC and facility performance are beneficial to advancing the practice of HPC design engineering and operations. At present, HPC facility managers at NERSC and elsewhere are experimentally verifying setpoints and component interactions, testing operational limits, and documenting their results. ODA, in combination with archived data, will allow these individual experiments to be shared, creating a growing body of experience for HPC and data center design. Typical rule-of-thumb approximations and historical practices will continue to be refined as ODA tools and processes become more widely adopted, further deepening industry experience.

ACKNOWLEDGMENTS

The authors would like to acknowledge the team from the Energy Efficient HPC Working Group Operational Data Analytics (ODA), with special thanks to Torsten Wilde, Ghaleb Abdulla, and Natalie

Bates for their technical leadership of the ODA Team and for spearheading the idea of this case study. Additional Berkeley Lab staff who are active contributors to the NERSC EE Project Team are Brent Draney, John Elliot, Mark Friedrich, Steve Greenberg, and Jingjing Liu. The NERSC OMNI System Development Team are Elizabeth Bautista, Cary Whitney, Thomas Davis, and Melissa Romanus of the Operations & Technology Group. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a US Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

REFERENCES

[1] French A. Avelar V., Azevedo D. 2014. PUE: A Comprehensive Examination of the Metric. (Feb. 2014). Retrieved May 7, 2019 from https://www.thegreengrid.org

[2] Davis T. Bautista E., Whitney C. 2016. Big Data Behind Big Data. In *Arora R. (eds)Conquering Big Data with High Performance Computing*. Springer, Cham, 163–189. https://doi.org/10.1007/978-3-319-33742-5_8

[3] Whitney C. Romanus M. Bourassa N. Bates N. Koenig G. Kubaska T. Bautista E., Davis T. 2019. Collecting, Monitoring, and Analyzing Facility and Systems Data at the National Energy Research Scientific Computing Center. In *Proceedings of the 48th International Conference on Parallel Processing (ICPP 2019)*, Vol. Energy Efficient HPC State of the Practice Workshop. ICPP 2019 Organizing Committee.

[4] Henry Coles. 2014. PUE in the Past; Towards TUE in the Future. (March 2014). Retrieved May 7, 2019 from https://eehpcwg.llnl.gov/assets/09_pue_past-tue_future.pdf