LIGHTHOUSE
WORLDWIDE SOLUTIONS

CASE STUDY

High Performance Computing (HPC) Facilities

National Energy Research Scientific Computing Center

Authored by:

Jason Kelly, Applications Director

Nicolas Ventura, Data Center Engineer

April, 2025

Introduction

The **National Energy Research Scientific Computing Cente**r (NERSC) is a premier high-performance computing facility dedicated to supporting scientific research. NERSC serves as the primary user facility for scientific computing for the U.S. Department of Energy’s Office of Science. Its mission is to provide high-performance computing and data resources to support a wide range of scientific research. NERSC is located at Lawrence Berkeley National Laboratory in Berkeley, California.

The Facility

An overview

More than 10,000 scientists from universities, national laboratories, and industry use NERSC's resources to conduct research in fields such as high energy physics, climate science, materials science, and more. NERSC houses state-of-the-art supercomputers, including the *Perlmutter* system, which is designed to handle large-scale scientific computations and data analysis.

NERSC supports a wide range of scientific projects, contributing to thousands of peer-reviewed publications each year. Its resources are crucial for advancing knowledge in areas like energy efficiency, environmental science, and fundamental physics.

NERSC has been a cornerstone in preserving scientific research data. Over the past 50 years, NERSC has utilized magnetic tape as a reliable and cost-effective medium to store vast amounts of scientific research. This method ensures the longevity and integrity of critical data, allowing researchers to access decades of valuable information. By maintaining controlled environmental conditions and periodically migrating data to newer formats, NERSC continues to safeguard the scientific legacy stored on magnetic tape, supporting ongoing and future research endeavors.

Background:

Operational Data Analytics (ODA) involves the continuous monitoring, archiving, and analysis of real-time performance data to provide actionable insights for operational decision-making. This approach is particularly valuable in complex environments like High-Performance Computing (HPC) systems, where it helps manage and optimize resources, improve efficiency, and ensure system reliability.

ODA typically encompasses several types of analytics:

- **Descriptive Analytics:** Summarizes past data to understand what has happened.
- **Diagnostic Analytics:** Examines data to determine why something happened.
- **Predictive Analytics:** Uses data to forecast future trends.
- **Prescriptive Analytics:** Provides recommendations based on data analysis.

By leveraging these analytics, organizations can make informed decisions, enhance operational performance, and address issues proactively.

OMNI

At NERSC, we put Operational Data Analytics into practice through the Operations Monitoring and Notification Infrastructure (OMNI) — a system designed to collect, analyze, and correlate operational data at scale. OMNI continuously ingests data from thousands of sensors, system metrics, and machine logs spanning the building management system, HPC resources, and supporting data center infrastructure into a centralized repository. On the front end, OMNI provides NERSC staff with real-time dashboards and alerting, historical data access, and analytic tools that enable both immediate monitoring and deeper insights over time.

Energy Saving Challenges

As a significant energy user of Lawrence Berkeley National Laboratory (Berkeley Lab), one goal of NERSC's ODA system is to maintain a certain level of energy efficiency. In general, HPC facilities face several challenges when it comes to energy savings:

High Energy Consumption:
HPC systems consume a significant amount of energy due to their intensive computational tasks. Managing this energy consumption is crucial for both cost and environmental reasons.

Cooling Requirements:
The heat generated by HPC systems requires efficient cooling solutions. Traditional cooling methods can be energy-intensive, adding to the overall energy consumption.

Resource Utilization:
Ensuring optimal utilization of computational resources is challenging. Underutilized resources can lead to wasted energy.

Scalability:
As HPC facilities scale up, managing energy efficiency becomes more complex. Larger systems require more sophisticated energy management strategies.

Data Center Infrastructure:
The infrastructure supporting HPC systems, including power distribution and cooling systems, must be designed for energy efficiency. Legacy systems may not be optimized for modern energy-saving techniques.

NERSC is tasked with striving for the best possible energy efficiency. Energy efficiency is quantified with PUE, or power usage effectiveness, which is calculated by the total, annualized energy divided by the annualized compute energy. A lower PUE describes a more energy-efficient facility, and a "perfect" PUE is 1.0. Sustainable Berkeley Lab's initiative is to maintain a PUE for NERSC at or below 1.1.

$$\text{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

NERSC is located in Building 59 on the Berkeley Lab campus. Berkeley is in the East Bay region of the San Francisco Bay Area. Due to its unique geography, Berkeley has a steady cool climate year-round, occasionally going above 80 degrees Fahrenheit in the late summer with only a handful of days in the low 90s. The highest wet bulb temperatures recorded by OMNI has been in the mid 70s. Due to these factors, the NERSC cooling facility does not rely on any chillers for additional cooling and rather relies solely on the climate for mechanical cooling.

The building's unique architecture takes full advantage of this. Located on the side of a hill, the west side of the mechanical floor is open to the atmosphere for direct air cooling and conditioning, and wind typically blows from the west. Cooled air is supplied into the raised floor in the data center above the mechanical space for air-cooled systems. On ideal days, NERSC uses single-pass air with evaporative cooling. A cooling water system with wet cooling towers are used for NERSC's liquid-cooled compute equipment.

The climate, geography, and architecture all contribute to significant energy savings for an energy-intensive facility such as NERSC. To maintain a cooling system with high performance, NERSC collects facility- and environmental-related metrics with OMNI, which provides data that helps in overcoming some of the challenges many HPC data centers face.

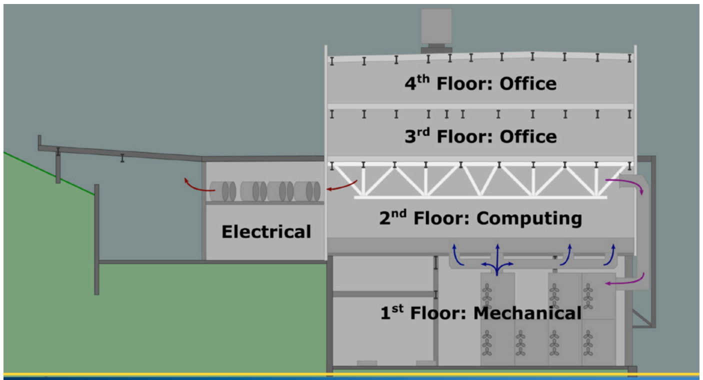

Building cross-section

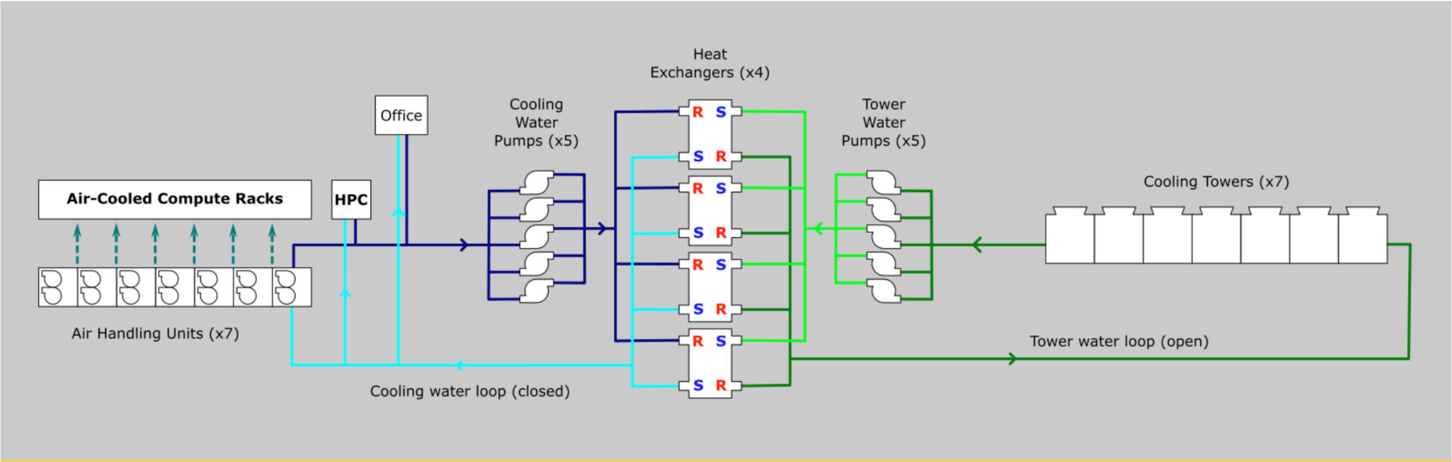

Building 59 Cooling Plant Systems

Air Quality

in Northern California

Despite all the natural benefits from NERSC's geography, it comes with significant challenges, particularly with air quality management. California experiences cool and wet winters when vegetation grows, which dries up during its hot and dry summers. Because of this, California has had an unfortunate history of wildfires, which cause severe destruction to its forests and land, and releases many particles into the atmosphere, which are wind-carried across the state.

In order to prolong the life of the tape archive and associated computer equipment, NERSC's HPC data center is held to the ISO 8 clean room standard. To ensure that the standard is met, several particle sensors were installed in various parts of the building for monitoring purposes. If particle counts are elevated for a certain length of time, the building infrastructure staff are instructed to enable CRAQ (computer room air quality) management mode. In this control mode, the data center is effectively isolated from outdoor air, and cooling is completely reliant on cooling water for air and liquid cooling. In this management mode, the building must be closed off for a certain length of time and can only return into normal mode if certain environmental conditions are met, so as to not allow moisture to condense on any electrical components. See "CRAQ Decision Flowchart" for more details.

CRAQ

Monitoring & Computer Room Air Quality

Incorporated into NERSC's building management system and operational strategy is CRAQ mode, which is designed to mitigate poor air quality from polluting the interior of the building and data center. This achieves the protection of HPC equipment, tape storage media, and human occupants. CRAQ mode is a binary operation, meaning that it's either enabled or disabled. When CRAQ mode is enabled, the air handling units' outside air dampers close to just 5%, and recirculate air, essentially isolating Building 59 from the external environment.

As previously stated, once CRAQ mode is enabled, Building 59 becomes completely reliant on its cooling water to maintain a stable temperature. Mechanical equipment will experience higher loads, thus leading to an increase in PUE. Liquid, air, and environmental conditions must be closely monitored during this time. At worst, the introduction of high temperatures or high particle counts within the data center may require the curtailment of compute resources to decrease the airflow in the data center.

When outdoor particle counts have lowered to an acceptable level, CRAQ mode can be disabled, given that certain criteria are met:

1. Particle counts outside must meet the ISO 8 clean room standard.
2. It must be less humid outside than inside.
3. The cooling water supply temperature must be greater than the outside air dew point temperature.

If (1) is not met, then introducing outside air could damage tape media equipment. If (2) or (3) are not met, then there is a risk of condensing moisture on the cooling pipes in the data center.

Due to these effects, enabling (or disabling) CRAQ mode is no simple decision. It requires extensive training and monitoring using the OMNI system. NERSC uses several Grafana dashboards in order to make decisions about CRAQ mode. Some of these dashboards include data from NERSC's particle sensors. Three particle sensors are located throughout the data center, three more are located within tape libraries, and one is located in the mechanical level, exposed to the outside air conditions. These sensors capture metrics on four different particle sizes; 0.3um, 0.5um, 1um, and 5um. The 0.5-micron particle sizes are most closely monitored, as these particles can most adversely affect the health of HPC and storage systems, whereas the larger sizes typically are better filtered out by NERSC's air handlers. Broad array of particle sensor data gives us the visibility needed to make CRAQ critical decisions in real time. It's also important to have a human in the loop in order to filter out false alarms, especially for the outdoor particle counts. For example, when semi-trucks drive by, they release a lot of pollution that is captured by our particle sensors.

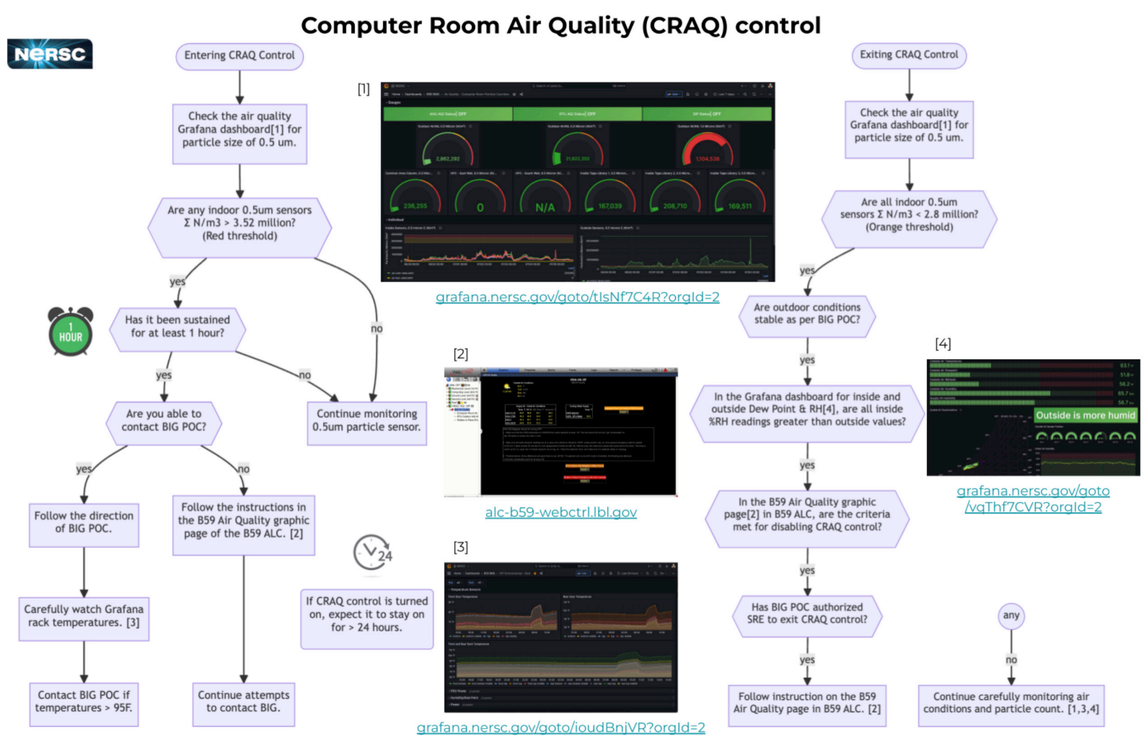

This flowchart is used by NERSC operators to maintain acceptable environmental air quality in the data center.

Dashboards

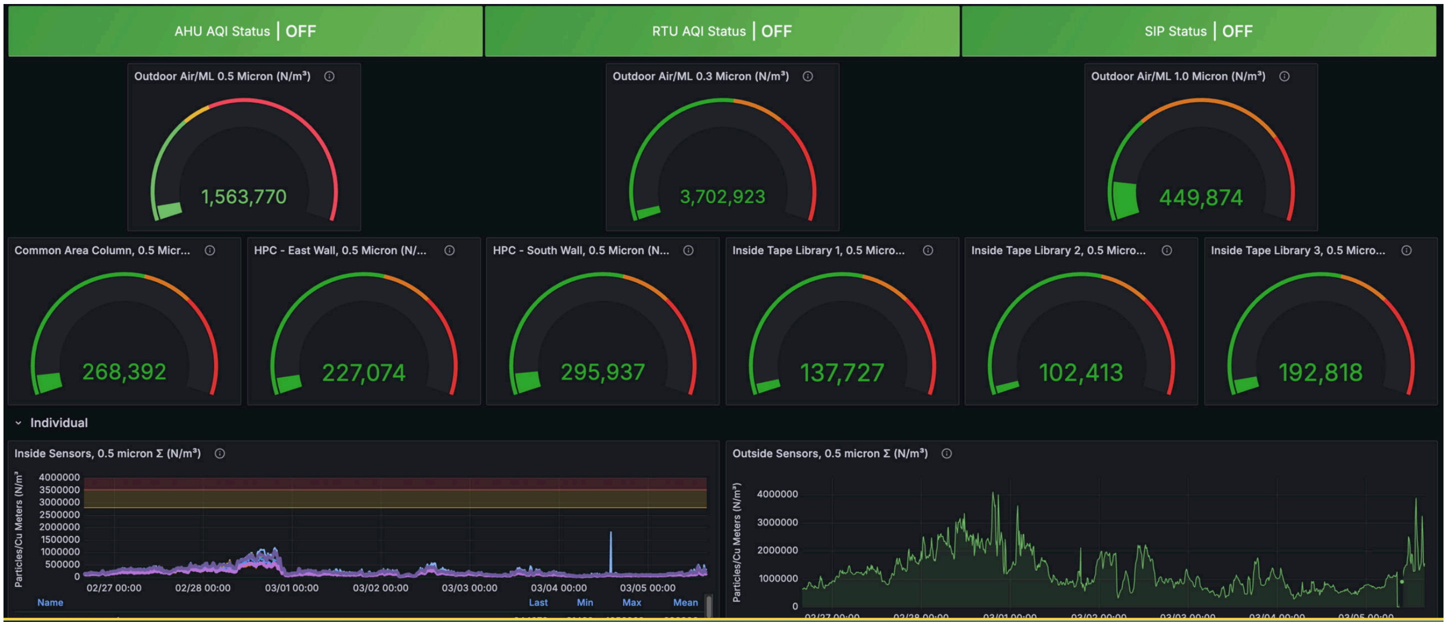

This is the primary air quality monitoring dashboard, which is always visible in one of the screens in the NERSC control room. This data is monitored from all of Building 59's air particle counters. We use the ISO 8 standards for our indoor alerting thresholds. The appropriate alerting thresholds for outdoor air conditions was about a x4.5 higher particle concentration than indoor air, which was determined from a data analysis.

This dashboard shows different qualities of outside air; particle count, psychrometrics, outside air damper positions, and "at-a-glance" panels which alert the operator of the risk of condensation.

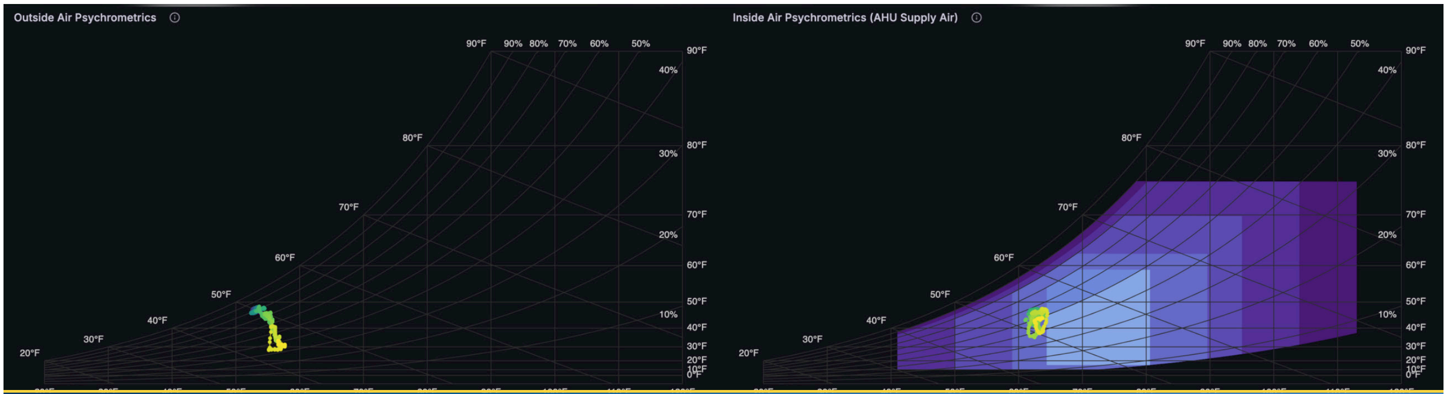

This panel shows two psychrometric charts. These compare outside and supply air psychrometrics. The purple envelopes in the supply air chart are defined by ASHRAE which outline environmental air conditions which must be met by data centers of various criticality levels. All the data shown in the dashboards above comes from our OMNI system.

System Overview

Particle counter technology

Apex R03p and ApexR3p sensors are deployed at several locations within the facility and also monitoring incoming air from outside. These sensors monitor particle sizes ranging from 0.3um to 5um with 0.1cfm and 1cfm flowrates. 0.3um and 0.5um particles are monitored and displayed.

In terms of particle sensitivity, to put into context, the average human hair is about 100um. Airborne particle counters use light scattering techniques where a particle is passed through a laser beam and it scatters light onto a photodetector. The size of the particle is relative to the scattered light that is detected by a device called a photodetector.

The photodetector converts light energy into an electrical amplitude where the amplitude is sized and the frequency of pulses from particles are also counted and recorded.

ApexRp Airborne Particle Counter with internal pump installed in the NERSC data center.

Particle Counter Sensor connectivity

The particle counters are connected over MODBUS TCP and connect into the OMNI system where critical data is recorded in real time and pulled from the MODBUS data string. Time of sample, sample volume, particle sizes and particle counts are collected and continuously trended on dashboards to provide real time data on the air quality outside the building and inside the ISO 8 data center where data archive servers, data cassettes, and tape, as well as the supercomputers, are located.

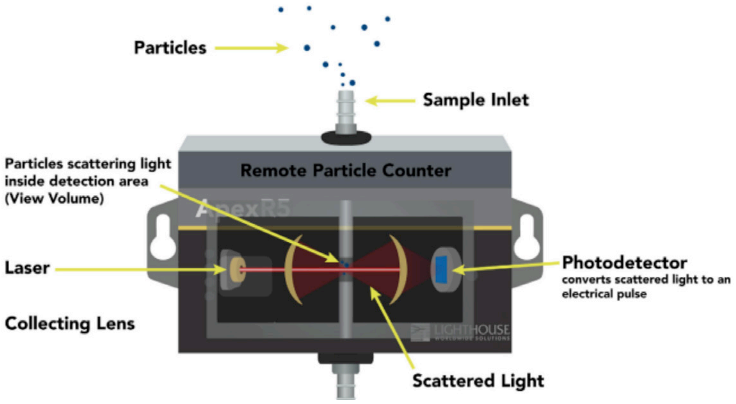

The Advantages

of Environmental Monitoring in HPC facilities

HPC data centers rely on precise environmental conditions to ensure optimal performance and longevity of critical hardware. Implementing particle counting environmental monitoring in these facilities offers several key advantages:

Improved Equipment Reliability & Longevity

Prevents Hardware Failures: Airborne particulate contamination can cause server overheating, circuit shorts, and physical damage to components like GPUs, CPUs, and power supplies.

Extends Component Lifespan: Reducing particulate matter minimizes wear and tear on sensitive microelectronics.

Optimized Cooling Efficiency

Reduces Airflow Blockages: Dust accumulation in cooling systems, heat sinks, and air filters impairs cooling efficiency, leading to increased energy consumption and equipment failures.

Enhances Air Filtration Effectiveness: Monitoring particulate levels ensures that air filtration systems (e.g., HEPA filters) are functioning properly and replaced when needed.

Energy and Cost Savings

Prevents Unplanned Downtime: Equipment failures due to particulate contamination can result in costly outages. Continuous monitoring allows for proactive maintenance.

Reduces Maintenance Costs: Identifying contamination sources early minimizes expensive emergency repairs and replacements.

Compliance with Industry Standards

Meets Cleanroom and ISO Standards: Some HPC data centers follow ISO 14644-1 cleanroom classifications to ensure air quality meets strict standards such as ISO 8 data centers.

Ensures Compliance in Regulated Environments: Industries such as financial services, healthcare, and scientific research often require strict environmental control to prevent data corruption or equipment failures.

Enhanced Operational Efficiency

Real-Time Monitoring for Rapid Response: Continuous particle monitoring provides immediate alerts when contamination levels exceed safe thresholds.

Predictive Maintenance: Data-driven insights help plan filter replacements and HVAC system maintenance before issues arise.

Protection Against External Contaminants

Minimizes Human-Introduced Contamination: Regular monitoring helps enforce protocols such as restricted access, improved airlocks, and proper attire to limit particulate ingress.

Detects External Pollution Sources: Construction work, exhaust fumes, and HVAC malfunctions can introduce contaminants that harm HPC equipment.

Better Risk Management and Business Continuity

Reduces Risk of Data Loss and Corruption: Prevents unexpected hardware failures that could lead to data loss in mission-critical applications.

Ensures High Availability and Uptime: Helps maintain 99.99% or higher uptime guarantees required by enterprise and cloud service providers.

It's all about the data in real time

It's all about the data in real time. In HPC data centers, where performance and uptime are critical, having immediate access to environmental data ensures that contamination risks are identified before they cause damage. Real-time particle monitoring prevents costly hardware failures by detecting airborne contaminants that can compromise cooling efficiency and damage sensitive components.

By continuously tracking air quality, operators can optimize filtration systems, reduce maintenance costs, and avoid unplanned downtime. Compliance with industry standards becomes seamless, as data-driven insights help maintain clean environments required for high-performance computing.

When every second counts, the ability to respond instantly to air quality changes means better risk management, enhanced operational efficiency, and uninterrupted service. Ensuring the integrity of an HPC facility isn't just about reacting to issues—it's about using real-time data to stay ahead of them.

Lessons Learned

The instrumentation in building 59 paired with OMNI system architecture has provided NERSC with a deeper understanding of real-time and long-term building operations and alerts. OMNI's model of "collect everything and then find what you need" has proven itself useful for energy and PUE reporting as well as preventative care for HPC systems such as alerting NERSC staff when to enable the CRAQ control mode using particle sensor data. It has also been a key player in appropriately sizing electrical equipment for power upgrades based on existing power data collection and analysis, which could lead to facility-related cost savings in the long run.

OMNI is mature in its data collection and visualization, but due to the thousands of metrics collected, creating actionable insights and dashboards is a long-term, iterative process. Although speculative, this could be augmented with the recent developments in artificial intelligence to notify operators in advance of anomalies or assist in troubleshooting issues based on data analysis, especially those which are dependent on metric types from several different sources.

Acknowledgments

This work was supported by the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory. The authors would like to thank Melissa Romanus[1], for her participation in the brainstorming and planning process, as well as her contributions about the OMNI system. The authors would also like to thank Norm Bourassa[2], Cary Whitney[3], and Thomas Davis[4] for their technical expertise.

[1] NERSC, OMNI Data Management Engineer

[2] NERSC, Building Infrastructure Group Lead (acting)

[3] NERSC, Operations Technology Group Lead (acting)

[4] NERSC, OMNI Lead Engineer